

Version 3.6

Part No. 311642-C Rev 00
April 2001

600 Technology Park Drive
Billerica, MA 01821-4130

Configuring the Contivity VPN Switch

NORTEL
NETWORKS™

Copyright © 2001 Nortel Networks

All rights reserved. April 2001.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel Networks Inc.

The software described in this document is furnished under a license agreement and may be used only in accordance with the terms of that license. The software license agreement is included in this document.

Trademarks

Nortel Networks, the Nortel Networks logo, Contivity, Preside, and Optivity are trademarks of Nortel Networks.

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

AXENT and AXENT OmiGuard Defender are trademarks of AXENT technologies, Inc.

Check Point FireWall VPN 1 and Firewall 1 are trademarks of Check Point Software Technologies, Ltd.

Cisco, Cisco 2514 Router, and Cisco Systems are trademarks of Cisco Technology, Inc.

Entrust and Entrust Authority are trademarks of Entrust Technologies, Inc.

Firewall-1 is a trademark of Check Point Software Technologies, Ltd.

Java and Java Script are trademarks of Sun Microsystems.

Linux and Linux FreeS/WAN are trademarks of Linus Torvalds.

Macintosh is a trademark of Apple Computer, Inc.

Microsoft, Active Directory, Windows, Hyperterminal, Microsoft Certification Authority, Microsoft Internet Explorer, and MS-DOS are trademarks of Microsoft Corporation.

NDS is a trademark of Novell Inc.

Netscape, Netscape Communicator, and Netscape Directory Server are trademarks of Netscape Communications Corporation.

Netware is a trademark of Novell, Inc.

NETVIEW is a trademark of International Business Machines Corp (IBM).

Novell and Novel intraNetWare are trademarks of Novell, Inc.

OPENVIEW is a trademark of Hewlett-Packard Company.

SafeNet/Soft-PK Security Policy Database Editor is a trademark of Information Resource Engineering, Inc. (IRE)

SecurID and Security Dynamics ACE Server are trademarks of RSA Security Inc.

SPECTRUM is a trademark of Cabletron Systems, Inc.

VeriSign and VeriSign OnSite are trademarks of VeriSign, Inc.

All other trademarks and registered trademarks are the property of their respective owners.

Restricted rights legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of conditions

In the interest of improving internal design, operational function, and/or reliability, Nortel Networks Inc. reserves the right to make changes to the products described in this document without notice.

Nortel Networks Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product may be Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

Nortel Networks Inc. software license agreement

NOTICE: Please carefully read this license agreement before copying or using the accompanying software or installing the hardware unit with pre-enabled software (each of which is referred to as “Software” in this Agreement). BY COPYING OR USING THE SOFTWARE, YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. THE TERMS EXPRESSED IN THIS AGREEMENT ARE THE ONLY TERMS UNDER WHICH NORTEL NETWORKS WILL PERMIT YOU TO USE THE SOFTWARE. If you do not accept these terms and conditions, return the product, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

1. License grant. Nortel Networks Inc. (“Nortel Networks”) grants the end user of the Software (“Licensee”) a personal, nonexclusive, nontransferable license: a) to use the Software either on a single computer or, if applicable, on a single authorized device identified by host ID, for which it was originally acquired; b) to copy the Software solely for backup purposes in support of authorized use of the Software; and c) to use and copy the associated user manual solely in support of authorized use of the Software by Licensee. This license applies to the Software only and does not extend to Nortel Networks Agent software or other Nortel Networks software products. Nortel Networks Agent software or other Nortel Networks software products are licensed for use under the terms of the applicable Nortel Networks Inc. Software License Agreement that accompanies such software and upon payment by the end user of the applicable license fees for such software.

2. Restrictions on use; reservation of rights. The Software and user manuals are protected under copyright laws. Nortel Networks and/or its licensors retain all title and ownership in both the Software and user manuals, including any revisions made by Nortel Networks or its licensors. The copyright notice must be reproduced and included with any copy of any portion of the Software or user manuals. Licensee may not modify, translate, decompile, disassemble, use for any competitive analysis, reverse engineer, distribute, or create derivative works from the Software or user manuals or any copy, in whole or in part. Except as expressly provided in this Agreement, Licensee may not copy or transfer the Software or user manuals, in whole or in part. The Software and user manuals embody Nortel Networks’ and its licensors’ confidential and proprietary intellectual property. Licensee shall not sublicense, assign, or otherwise disclose

to any third party the Software, or any information about the operation, design, performance, or implementation of the Software and user manuals that is confidential to Nortel Networks and its licensors; however, Licensee may grant permission to its consultants, subcontractors, and agents to use the Software at Licensee's facility, provided they have agreed to use the Software only in accordance with the terms of this license.

3. Limited warranty. Nortel Networks warrants each item of Software, as delivered by Nortel Networks and properly installed and operated on Nortel Networks hardware or other equipment it is originally licensed for, to function substantially as described in its accompanying user manual during its warranty period, which begins on the date Software is first shipped to Licensee. If any item of Software fails to so function during its warranty period, as the sole remedy Nortel Networks will at its discretion provide a suitable fix, patch, or workaround for the problem that may be included in a future Software release. Nortel Networks further warrants to Licensee that the media on which the Software is provided will be free from defects in materials and workmanship under normal use for a period of 90 days from the date Software is first shipped to Licensee. Nortel Networks will replace defective media at no charge if it is returned to Nortel Networks during the warranty period along with proof of the date of shipment. This warranty does not apply if the media has been damaged as a result of accident, misuse, or abuse. The Licensee assumes all responsibility for selection of the Software to achieve Licensee's intended results and for the installation, use, and results obtained from the Software. Nortel Networks does not warrant a) that the functions contained in the software will meet the Licensee's requirements, b) that the Software will operate in the hardware or software combinations that the Licensee may select, c) that the operation of the Software will be uninterrupted or error free, or d) that all defects in the operation of the Software will be corrected. Nortel Networks is not obligated to remedy any Software defect that cannot be reproduced with the latest Software release. These warranties do not apply to the Software if it has been (i) altered, except by Nortel Networks or in accordance with its instructions; (ii) used in conjunction with another vendor's product, resulting in the defect; or (iii) damaged by improper environment, abuse, misuse, accident, or negligence. **THE FOREGOING WARRANTIES AND LIMITATIONS ARE EXCLUSIVE REMEDIES AND ARE IN LIEU OF ALL OTHER WARRANTIES EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.** Licensee is responsible for the security of its own data and information and for maintaining adequate procedures apart from the Software to reconstruct lost or altered files, data, or programs.

4. Limitation of liability. IN NO EVENT WILL NORTEL NETWORKS OR ITS LICENSORS BE LIABLE FOR ANY COST OF SUBSTITUTE PROCUREMENT; SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES; OR ANY DAMAGES RESULTING FROM INACCURATE OR LOST DATA OR LOSS OF USE OR PROFITS ARISING OUT OF OR IN CONNECTION WITH THE PERFORMANCE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF NORTEL NETWORKS RELATING TO THE SOFTWARE OR THIS AGREEMENT EXCEED THE PRICE PAID TO NORTEL NETWORKS FOR THE SOFTWARE LICENSE.

5. Government licensees. This provision applies to all Software and documentation acquired directly or indirectly by or on behalf of the United States Government. The Software and documentation are commercial products, licensed on the open market at market prices, and were developed entirely at private expense and without the use of any U.S. Government funds. The license to the U.S. Government is granted only with restricted rights, and use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1) of the Commercial Computer Software—Restricted Rights clause of FAR 52.227-19 and the limitations set out in this license for civilian agencies, and subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause of DFARS 252.227-7013, for agencies of the Department of Defense or their successors, whichever is applicable.

6. Use of software in the European Community. This provision applies to all Software acquired for use within the European Community. If Licensee uses the Software within a country in the European Community, the Software Directive enacted by the Council of European Communities Directive dated 14 May, 1991, will apply to the examination of the Software to facilitate interoperability. Licensee agrees to notify Nortel Networks of any such intended examination of the Software and may procure support and assistance from Nortel Networks.

7. Term and termination. This license is effective until terminated; however, all of the restrictions with respect to Nortel Networks' copyright in the Software and user manuals will cease being effective at the date of expiration of the Nortel Networks copyright; those restrictions relating to use and disclosure of Nortel Networks' confidential information shall continue in effect. Licensee may terminate this license at any time. The license will automatically terminate if

Licensee fails to comply with any of the terms and conditions of the license. Upon termination for any reason, Licensee will immediately destroy or return to Nortel Networks the Software, user manuals, and all copies. Nortel Networks is not liable to Licensee for damages in any form solely by reason of the termination of this license.

8. Export and re-export. Licensee agrees not to export, directly or indirectly, the Software or related technical data or information without first obtaining any required export licenses or other governmental approvals. Without limiting the foregoing, Licensee, on behalf of itself and its subsidiaries and affiliates, agrees that it will not, without first obtaining all export licenses and approvals required by the U.S. Government: (i) export, re-export, transfer, or divert any such Software or technical data, or any direct product thereof, to any country to which such exports or re-exports are restricted or embargoed under United States export control laws and regulations, or to any national or resident of such restricted or embargoed countries; or (ii) provide the Software or related technical data or information to any military end user or for any military end use, including the design, development, or production of any chemical, nuclear, or biological weapons.

9. General. If any provision of this Agreement is held to be invalid or unenforceable by a court of competent jurisdiction, the remainder of the provisions of this Agreement shall remain in full force and effect. This Agreement will be governed by the laws of the state of California.

Should you have any questions concerning this Agreement, contact Nortel Networks Inc., 2375 N. Glenville Dr., Richardson, TX 75082.

LICENSEE ACKNOWLEDGES THAT LICENSEE HAS READ THIS AGREEMENT, UNDERSTANDS IT, AND AGREES TO BE BOUND BY ITS TERMS AND CONDITIONS. LICENSEE FURTHER AGREES THAT THIS AGREEMENT IS THE ENTIRE AND EXCLUSIVE AGREEMENT BETWEEN NORTEL NETWORKS AND LICENSEE, WHICH SUPERSEDES ALL PRIOR ORAL AND WRITTEN AGREEMENTS AND COMMUNICATIONS BETWEEN THE PARTIES PERTAINING TO THE SUBJECT MATTER OF THIS AGREEMENT. NO DIFFERENT OR ADDITIONAL TERMS WILL BE ENFORCEABLE AGAINST NORTEL NETWORKS UNLESS NORTEL NETWORKS GIVES ITS EXPRESS WRITTEN CONSENT, INCLUDING AN EXPRESS WAIVER OF THE TERMS OF THIS AGREEMENT.

Contents

Preface	23
Before you begin	23
Text conventions	23
Related publications	25
Acronyms	26
How to get help	27
 Chapter 1	
Overview	29
The Contivity VPN Switch	29
Extranet access	30
Remote access versus extranet access	31
Network configurations	32
Switch behind the router, parallel with firewall	33
Switch with direct connections to the PDN	33
Switch behind a firewall and a router	34
Branch office configuration	35
Routing	36
Routing policy service (RPS)	37
Routing table	37
About RIP	38
About OSPF	38
About VRRP	38
Integrated firewall solutions	39
Private and public interfaces	40
IP addressing	41
Tunnels	44
IPSec	45
PPTP	46
L2TP	47
L2TP over IPSEC branch office support	47
L2F	47

Branch office connections	48
Network Address Translation (NAT)	49
Authentication	49
LDAP	50
RADIUS	51
SSL and digital certificates	51
Tunnel certificates	52
Encryption	52
Filters for access control	53
Management	53
SNMP traps	54
Quality of service	54
Differentiated Services (DiffServ)	55
Bandwidth management	55
Forwarding priority	56
Call admission priority	57
RSVP	58
Accounting	59
Local configuration and LDAP file structures	59
LDAP database	59
Configuration file	61
Accounting and logging	62
Command line interface	62
Serial point-to-point protocol (PPP)	63
FIPS	63

Chapter 2

Managing an Enterprise or Carrier Environment 65

Network configurations	66
Accessing the Contivity VPN Switch	68
Configuration tasks	69
Split tunneling	72
Safe mode configuration	73
Restricted mode	73
Nailed-up control tunnels	74

Setting up branch office tunnels	74
Using bulkload scripts	74
Working with user groups	74
Setting up servers	75
Monitoring the switch	75
Configuration file management	76
System logging	76
SNMP support	76
SNMP traps	77
Trap on hardware warnings and alerts	77
Routing table	77
Management tasks	78
Switch health	79
Switch data storage	79
RADIUS accounting	80
Accounting data	80
Accounting records	81
Service level agreements	82
Backup/restore	83
Upgrades	83
Management applications	83
VPN manager	83
Preside	84
Chapter 3	
Configuring the Switch	87
Licensing	88
IP Address Configuration Utility	88
Serial Interface configuration	91
Safe mode configuration	92
Reconfiguring IP address values	93
Changing management IP address when interface IP address is defined	94
Running the IP address configuration utility	94
Startup configuration requirements	95
Management IP address	96

Subnet mask	96
Gateway IP address	96
Web browser	97
Login and password	97
Preparing for configuration	97
Welcome screen	98
LAN interfaces	102
WAN interfaces	103
T-1 with integrated CSU/DSU	103
Configuring the T-1 interface with an integrated CSU/DSU	104
Verification	105
Alarm conditions	105
LEDs	106
Troubleshooting	106
Limitations	106
Hardware encryption accelerator	107
Hardware encryption accelerator and load balancing	109
Hardware encryption accelerator security aspects	109
Network Time Protocol (NTP)	109

Chapter 4

Configuring tunnels..... 111

Configuring user tunnels	116
Configuring branch offices	119
Branch-to-branch with a firewall and a router	120
Branch-to-branch through a gateway switch	121
Using NAT	122
Translation types	124
Overlapping branch office addresses	127
Creating NAT sets	128
Sample branch office configuration procedure	129
Configuring the local switch	132
Configuring the remote switch	138
Configuring the remote branch office connection	139
Control tunnels	140

Control tunnel types	142
Restricted mode	143
Nailed-up control tunnels	144
Creating control tunnels	144
Creating a user control tunnel from the serial interface	147
IPSec client features	148
Split tunneling	149
Third-party IPSec clients	150
Forced logout	152
Client fail-over	152
Client auto connect	153
Banner	154
Password storage	154
Client screen saver	154
Domain name	154
Client policy	155
Configuring L2TP over IPSec	155
Switch configuration	155
Windows 2000 configuration	160
Branch office	161
Proxy ARP and tunnel to tunnel traffic	163

Chapter 5

Configuring Routing, Firewalls, and IPX 165

Configuring routing	165
Static routes	167
VPN routing	167
Enhanced routing	168
Services routing	169
Routing policy service (RPS)	170
Routing table	173
Route lookup	175
Default routes	176
Rules of redistribution for RIP and OSPF	176
About RIP	177

Using RIP on your switch	178
Protecting against routing loops	178
About OSPF	179
Using OSPF	179
Advanced routing and firewall keys	180
About VRRP	180
Configuring VRRP	181
Traffic patterns	181
Client address redistribution	183
Configuring firewalls	186
Contivity Stateful Firewall	186
Contivity Tunnel Filter	187
Check Point FireWall-1	187
Enabling your switch's firewall support	187
Confirming and saving your firewall selection	188
Configuring IPX	188
IPX client	190
Windows 95 and Windows 98	190
Windows NT	190
IPX group configuration	190
Sample IPX VPN switch topology	190
Chapter 6	
Configuring Servers	193
Authentication services	193
LDAP, RADIUS, and the IPsec client	197
LDAP database authentication	198
RADIUS-based authentication	198
RADIUS authentication servers	200
RADIUS authentication class attribute values	200
RADIUS-Assigned Framed-IP-Address attribute	202
Configuring the switch for RADIUS	202
IPsec and RADIUS	202
PPTP and RADIUS	203
AXENT software tokens and the IPsec client	204

LDAP servers	204
Using a Netscape directory server	206
Quick start	206
Advanced configuration	208
Schema/attribute index cleanup	208
Schema/attribute index modifications	209
SSL	209
Exporting internal LDAP to Netscape directory server	211
RADIUS accounting configuration	213
Remote user IP address pool	213
Configuring the switch as a RADIUS server	214
 Chapter 7	
Using Certificates	215
 SSL and digital certificates	215
Tunnel certificates	218
Public Key Infrastructure (PKI)	220
CA and X.509 certificates	220
Certificate application, creation, and installation	220
Entrust	220
VeriSign	221
Certificate lifetime	221
Certificate use	222
Certificate maintenance, expiration, and termination	222
Certification authority network controls	223
Control of public and private keys	223
Certificate revocation list (CRL)	225
CRL servers	226
Using digital certificates	226
Getting a server certificate	227
Generating a server certificate request	227
Requesting a server certificate	228
Obtaining a server certificate from the CA	229
Importing the server certificate to the switch	230
Obtaining the CA certificate from the PKI	231

Importing the CA certificate to the switch	231
Enabling RSA digital signature IPSec authentication	232
Verifying your time zone setting	233
Adding remote clients to the PKI	233
Entrust server	235
Contivity VPN Client (Entrust)	235
Contivity VPN Client (VeriSign)	235
Authenticating a client with a digital certificate	236
Certificate maintenance	237
Microsoft certificate services	237

Chapter 8

Configuring for Interoperability 239

Configuring the Cisco 2514 router, version 11.3	239
Configuring the switch for Cisco interoperability	242
Configuring the Check Point VPN-1 and FireWall-1, Version 4.0	243
Establishing an encrypted tunnel on the switch side	244
Establishing an encrypted tunnel on the firewall side	245
Establishing a tunnel from the switch to the firewall	246
Establishing a tunnel to and from the firewall	247
Configuring the firewall's public interface	248
Configuring the switch authentication schemes	249
Configuring the firewall's encryption methods	250
Configuring the switch's public interface	252
Configuring the switch's encryption methods	253
Adding rules to the configurations	255
Editing encryption for rules	256
Configuring the switch for Check Point interoperability	259
Configuring the SafeNet/Soft-PK Security Policy Database Editor, version 1.0s ...	260
Connecting to IRE SafeNET/Soft-PK Security Policy Client	262
Configuring the switch for IRE interoperability	266
Third-party client installation	266
Configuring the switch as a branch office tunnel	267
Configuring the switch as a user tunnel	268

Chapter 9	
Switch Administration	271
Administrators	271
Backups	272
Tools	274
Recovery	275
Using the recovery diskette	275
Upgrades	279
System configuration	280
File management	280
SNMP	281
Shutdown	281
Switch status	282
Sessions	284
Reports	285
System	285
Health check	285
Statistics	286
Firewall	286
Accounting	286
Event log	286
System log	287
Security log	287
Configuration log	287
Chapter 10	
Troubleshooting	289
Troubleshooting tools	290
Client-based tools	290
Switch-based tools	291
Other tools	292
Solving connectivity problems	292
Diagnosing client connectivity problems	293
Common client connectivity problems	294
Problems with name resolution using DNS services	297

Network browsing problems	298
Diagnosing WAN link problems	300
Hardware encryption accelerator connectivity	302
Solving performance problems	303
Eliminating modem errors	303
Using DES encryption instead of Triple DES when using IPSec	303
Hardware encryption accelerator performance	304
Additional information	313
Solving general problems	314
Web browser problems and the Contivity VPN Client Manager	314
Enabling Web browser options	314
Web browser error messages	316
Reporting a problem with a Web browser	318
Switch problems	318
Solving routing problems	320
Client address redistribution problems	320

Appendix A

MIB Support 323

SNMP RFC support	323
Novell IPX MIB	323
Novell RIP-SAP MIB	323
RFC 1850 -- OSPF Version 2 Management Information Base	324
RFC 1724 -- RIP Version 2 MIB Extension	324
RFC 1213 -- Network Management of TCP/IP-Based Internets MIB	324
RFC 2667 -- IP Tunnel MIB	324
RFC 2787 -- VRRP MIB	324
RFC 2737 -- Entity MIB	325
RFC 1573 -- IanaIfType MIB	326
RFC 2233 -- If MIB	326
RFC 2571-- Snmp-Framework MIB	326
CES MIB	326
cestraps.mib -- Nortel Networks proprietary MIB	327
newoak.mib	328
Hardware-related traps	329

Server-related traps	333
Software-related traps	335
Login-related traps	335
Intrusion-related traps	336
System-related traps	336
Information passed with every trap	337
 Appendix B	
Using Serial PPP	341
Establishing a serial PPP connection	341
Setting up a Dial-Up Networking connection	342
Setting up the modem	343
Setting up the switch	343
Dialing in to the switch	345
Troubleshooting Serial PPP	345
PPP option settings	347
 Appendix C	
System Messages	349
Certificate messages	349
ISAKMP messages	351
Branch office messages	354
SSL messages	355
Database messages	356
Security messages	357
RADIUS accounting messages	367
RADIUS authentication messages	370
Routing messages	374
Hardware messages	380
Glossary	383
 Index	391

Figures

Figure 1	The Internet and an intranet form an extranet	31
Figure 2	Traditional remote access environment	32
Figure 3	Direct connections to a router	33
Figure 4	Independent WAN connections	34
Figure 5	Switch located behind a firewall	34
Figure 6	Typical branch office environment	36
Figure 7	Private and public interfaces	41
Figure 8	Sample IP addressing scheme	42
Figure 9	Sample IPX addressing scheme	43
Figure 10	Remote user to PDN tunnel and on to the switch	45
Figure 11	Typical branch office environment	48
Figure 12	Switch sample NAT environment	49
Figure 13	ISP network configuration	67
Figure 14	Branch office enterprise network configuration	68
Figure 15	Welcome screen	99
Figure 16	Navigational menu	101
Figure 17	LAN-to-switch connection	102
Figure 18	Switch-to-PDN configuration	103
Figure 19	Tunnel connection configuration	111
Figure 20	Base, parent, and child group associations	114
Figure 21	Typical branch office environment	120
Figure 22	Branch-to-branch with a firewall and a router	121
Figure 23	Indirectly connected branch offices	122
Figure 24	Sample NAT environment	123
Figure 25	Static address translation	125
Figure 26	Port address translation	126
Figure 27	Pooled address translation	127
Figure 28	Overlapping address translation	128
Figure 29	Sample branch office configuration	130

Figure 30	Setting up a branch office configuration	131
Figure 31	Branch office control tunnel	141
Figure 32	Sample control tunnel environment	142
Figure 33	Sample split tunneling environment	149
Figure 34	Traffic flow	166
Figure 35	VPN routing: allowed traffic patterns	168
Figure 36	Enhanced routing: allowed traffic patterns	169
Figure 37	Router configuration with OSPF and RIP	171
Figure 38	Accept and announce policies	172
Figure 39	Switch routing table	175
Figure 40	Client address redistribution	184
Figure 41	Summarization for Client Address Redistribution	185
Figure 42	IPX topology	191
Figure 43	Contivity VPN Switch and authentication servers	194
Figure 44	Authentication server validation flowchart	196
Figure 45	IPSec client authentication options	197
Figure 46	RADIUS authentication class attribute values	201
Figure 47	Sample certificate chain	217
Figure 48	Digital certificate components	219
Figure 49	Obtaining a server certificate	227
Figure 50	Enabling server-side CRL usage	232
Figure 51	Certificate enrollment and maintenance	234
Figure 52	Switch and Cisco 2514 network topology	240
Figure 53	Switch and Check Point FireWall-1 network	243
Figure 54	Switch and IRE SafeNet network topology	261
Figure 55	Split tunneling example	269
Figure 56	Backup servers 1 and 2	273
Figure 57	The Recovery Diskette screen	276
Figure 58	Nortel Networks logging scheme	284

Tables

Table 1	Firewall filter pass-through ports	35
Table 2	Sample IP addressing associations	42
Table 3	Sample IPX addressing associations	43
Table 4	Bandwidth allocation per priority level	56
Table 5	Call admission priority	57
Table 6	Maximum connections per priority	58
Table 7	Configuration tasks	71
Table 8	Management tasks	78
Table 9	Field IDs for data collection records	80
Table 10	Configuration options	87
Table 11	Web interface configuration options	98
Table 12	Alarm conditions and LEDs	105
Table 13	Mapping minimum data protection levels to encryption levels	158
Table 14	User group properties used by IPSec transport connections	159
Table 15	Redistribution rules	176
Table 16	Traffic patterns	182
Table 17	RADIUS example details	201
Table 18	Troubleshooting tools	292
Table 19	Trap categories	338
Table 20	DIP switches configuration	343

Preface

This guide describes the Nortel Networks* Contivity* VPN Switch. It also provides complete details to help you configure, monitor, and troubleshoot your Contivity VPN Switch.

Before you begin

This guide is for network managers who are responsible for setting up client software for the Contivity VPN Switch. This guide assumes that you have the following background:

- Experience with windowing systems or graphical user interfaces (GUIs)
- Familiarity with the network management

This guide refers to the Contivity VPN Switch as the switch.

Text conventions

This guide uses the following text conventions:

angle brackets (< >) Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command.

Example: If the command syntax is
`ping <ip_address>`, you enter
`ping 192.32.10.12`

bold Courier text Indicates command names and options and text that you need to enter.

Example: Use the **ping** command.

braces ({})	Indicate required elements in syntax descriptions where there is more than one option. You must choose only one of the options. Do not type the braces when entering the command. Example: If the command syntax is <code>show ip {alerts routes}</code>, enter either <code>show ip alerts</code> or <code>show ip routes</code>, but not both.
brackets ([])	Indicate optional elements in syntax descriptions. Do not type the brackets when entering the command. Example: If the command syntax is <code>show ip interfaces [-alerts]</code>, you can enter either <code>show ip interfaces</code> or <code>show ip interfaces -alerts</code>.
ellipsis points (. . .)	Indicate that you repeat the last element of the command as needed. Example: If the command syntax is <code>ethernet/2/1 [<parameter> <value>] . . .</code>, you enter <code>ethernet/2/1</code> and as many parameter-value pairs as needed.
<i>italic text</i>	Indicates new terms, book titles, and variables in command syntax descriptions. Where a variable is two or more words, the words are connected by an underscore. Example: If the command syntax is <code>show at <valid_route></code>, <i>valid_route</i> is one variable and you substitute one value for it.
plain Courier text	Indicates command syntax and system output, for example, prompts and system messages. Example: <code>Set Trap Monitor Filters</code>
arrow (→)	Shows menu paths. Example: <code>Protocols→IP</code> identifies the IP option on the Protocols menu.

vertical line ()	Separates choices for command keywords and arguments. Enter only one of the choices. Do not type the vertical line when entering the command. Example: If the command syntax is <code>show ip {alerts routes}</code>, you enter either <code>show ip alerts</code> or <code>show ip routes</code>, but not both.
asterisk (*)	Indicates a trademark. See the Title Page for trademark information. Example: Contivity* Nortel Networks, the Nortel Networks logo, and Contivity are trademarks of Nortel Networks.

Related publications

For more information about the Contivity VPN Switch, refer to the following publications:

- Release notes for the switch (part number 301459-T) and the client (part number 301459-V) provide the latest information, including brief descriptions of the new features, problems fixed in the this release, and known problems and workarounds.
- *Reference for the Contivity VPN Switch* (part number 311643-C) describes details for the fields on the User Interface screens.
- *Reference for the Contivity VPN Switch Command Line Interface* (part number 311645-B) describes the commands that you can use from the command line interface.
- *Managing the Contivity Stateful Firewall* (part number 312538-A) describes firewall concepts, how to configure and monitor the firewall, and the firewall commands that you can use from the command line interface.
- *Installing the Extranet Access Client* (part number 311644-A) provides procedural information to help you configure, monitor, and troubleshoot your switch.

You can print selected technical manuals and release notes free, directly from the Internet. Go to the www.nortelnetworks.com/documentation URL. Find the product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Use Adobe Acrobat Reader to open the manuals and release notes, search for the sections you need, and print them on most standard printers. Go to Adobe* at the www.adobe.com URL to download a free copy of the Adobe Acrobat Reader*.

You can purchase selected documentation sets, CDs, and technical publications through the Internet at the www1.fatbrain.com/documentation/nortel/ URL.

Acronyms

This guide uses the following acronyms:

ACK	acknowledgement
CA	certificate authority
CHAP	Challenge Handshake Authentication protocol
CRL	certificate revocation list
DN	distinguished name
DNS	Domain name system
EAC	Extranet Access Client
FIPS	Federal Information Processing Standards
FTP	File Transfer Protocol
IP	Internet Protocol
ISAKMP	Internet Security Association and Key Management Protocol
IKE	IPSec Key Exchange
L2TP	Layer2 Tunneling Protocol
LDAP	Lightweight Directory Access Protocol
LAN	local area network
MAC	media access control address

NAT	network address translation
NTP	Network Time Protocol
OSPF	Open Shortest Path First
PPP	Point-to-Point Protocol
PPTP	Point-to-Point Tunneling Protocol
RSVP	Resource Reservation Protocol
RIP	Routing Information Protocol
SNMP	Simple Network Management Protocol
UDP	User Datagram Protocol
URL	uniform resource locator
VRRP	Virtual Router Redundancy Protocol
WAN	wide area network

How to get help

If you purchased a service contract for your Nortel Networks product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Nortel Networks service program, contact one of the following Nortel Networks Technical Solutions Centers:

Technical Solutions Center	Telephone
EMEA	(33) (4) 92-966-968
North America	(800) 2LANWAN or (800) 252-6926
Asia Pacific	(61) (2) 9927-8800
China	(800) 810-5000

An Express Routing Code (ERC) is available for many Nortel Networks products and services. When you use an ERC, your call is routed to a technical support person who specializes in supporting that product or service. To locate an ERC for your product or service, go to the www12.nortelnetworks.com/ URL and click ERC at the bottom of the page.

Chapter 1

Overview

This chapter introduces the Nortel Networks Contivity VPN Switch. The switch provides scalable, secure, and manageable remote access using the Internet.

The switch combines remote access protocols, security, authentication, authorization, and encryption technologies into a single solution. The switch can support up to 5000 simultaneous encrypted tunnels (depending on the model). An individual user or group of users can be associated with a set of attributes that provide custom access to an extranet. In effect, you can create a personal extranet based on the special needs of a user or group.

The Contivity VPN Switch

The switch includes the most popular tunneling protocols, IP Security (IPSec), Point-to-Point Tunneling Protocol (PPTP), Layer 2 Forwarding Tunneling Protocol (L2TP), and Layer 2 Forwarding (L2F). IPSec uses digital certificates, password-based keys, and tokens for authentication; PPTP, L2TP, and L2F use Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP) for authentication. The PPTP implementation for the switch supports MS-CHAP authentication with 56- to 128-bit key encryption.

The switch provides more security than traditional remote access schemes by using a combination of authorization, authentication, privacy, and access control for each user. In addition, support for the IPSec protocol and related Internet Security Association and Key Management Protocol (ISAKMP) and the Oakley Key Establishment Protocol further enhances security.

For authentication and access control, the switch supports an internal or external Lightweight Directory Access Protocol (LDAP) server and external Remote Authentication Dial-In User Service (RADIUS) servers.

To restrict access, the switch uses packet filtering based on protocol ID, direction, source and destination IP addresses, source and destination ports, and TCP connection establishment. Additionally, Nortel Networks provides a set of predefined filters that you can use either directly or tailor to your network needs.

Call admission and packet-forwarding priorities, and support for Resource ReSerVation Protocol (RSVP) provide unique quality of service methods.


The HTML and Java* Web management interface allows different switch administrators to have different rights, including configuration, status, and monitoring. The switch offers RADIUS accounting support and extensive logging, including events, system, configuration, and security logs.

The 4000 series of the Contivity VPN Switch is fully redundant, featuring dual 404-watt fault tolerant power supplies, dual-mirrored hard drives, and automatic backup server support. Other models have a single power supply and hard drive.

Extranet access


With its combination of secure, manageable, and scalable features, the switch offers significant extranet access savings. You can shift information technology resources from solving the current remote user access problems to other, more proactive administrative and management areas. And you can eliminate modem-management pool problems from your organization and shift them to your extranet provider.

Extranet access allows remote users to dial in to an Internet Service Provider (ISP) anywhere in the world and reach corporate headquarters or branch offices using the ISP. The extranet provides remote users access to corporate databases, mail servers, and file servers. [Figure 1](#) shows a typical extranet environment.

Figure 1 The Internet and an intranet form an extranet

Remote access versus extranet access

Traditional corporate remote access environments employ banks of modems to handle incoming service requests. The switch allows Internet Service Providers (ISPs) to take over the role of point-of-presence (POP) providers of modem access. This frees corporate personnel to perform other proactive duties. The switch improves performance while lowering overhead. And, the resource shift translates to significant corporate savings. [Figure 2](#) shows a remote access environment.

Figure 2 Traditional remote access environment

You can also save by reducing the number of telephone company (Telco) T1 lines needed. Traditional T1 (channelized) lines support only 24 connections per line; a T1 connected to a switch aggregates many sessions through statistical multiplexing and can service between 100 and 200 extranet users, depending upon their usage.

Additionally, there are savings from the reduced Telco expenses for the long-distance calls from employee homes to corporate modem pool access points, as well as replacement of expensive private leased lines that link remote offices. Extranet users save by dialing into the local POP, which is usually a free call.

Network configurations


The switch can be integrated into your network in several configurations, depending on your network topology and method of Internet access. Four common topologies are described in this section:

- Switch behind the router, in parallel with a firewall
- Switch with direct connections to the Public Data Network (PDN)
- Switch behind a firewall and a router
- Branch office configuration

Switch behind the router, parallel with firewall


This configuration shows the switch connected behind a router and in parallel with a firewall. This configuration allows you to retain your preexisting configuration while still allowing secure extranet access through the switch. One of the major benefits of this configuration is that it shares an existing connection to the PDN. [Figure 3](#) shows a firewall and the switch with connections to a router.

Figure 3 Direct connections to a router




Switch with direct connections to the PDN

In this configuration, the firewall and router are completely independent of the switch. Their connections and configuration remain the same. The switch is installed like a traditional remote access server except that the wide area network (WAN) connection is to the PDN rather than to a dial-up telephone network. This allows extranet bandwidth to grow without impacting other Internet traffic. [Figure 4](#) shows the switch with independent, direct connections to a WAN.

Figure 4 Independent WAN connections

Switch behind a firewall and a router

This configuration shows the switch connected behind both a firewall and a router. For this configuration to work, the router and firewall must be configured to allow tunnels to pass through to the switch. This is usually a simple matter that in most cases requires changes only to the firewall. [Figure 5](#) shows the switch behind a firewall and a router.

Figure 5 Switch located behind a firewall


In the preceding configuration, you have to modify your firewall filters to pass PPTP, L2F, and other tunneled traffic through to the switch. Refer to Table 1 for details.

Table 1 Firewall filter pass-through ports

Service	Protocol number	Source port	Destination port
PPTP Control Connection	6 (TCP)	>1023	1723
PPTP Tunnel Encapsulation	47 (GRE)	N/A	N/A
ISAKMP/ IPSec Key Management	17 (UDP)	500	500
IPSec Tunnel Encapsulation	50 (ESP) 51 (AH)	N/A	N/A
L2TP/L2F	17 (UDP)	>1023	1701
Reservation Control	46 (RSVP)	N/A	N/A

Branch office configuration

The branch office feature allows you to configure an IPSec tunnel connection between two private networks. Typically, one private network is behind a locally configured switch while the other is behind a remote switch. Branch office configuration allows you to configure the accessible subnetworks behind each switch. The configuration also contains the information that is necessary to set up the connection, such as the switch's IP addresses, encryption types and authentication methods. Local policy restrictions such as access hours, filter sets, and call admission priorities can be applied to limit connectivity into local subnetworks. [Figure 6](#) shows a typical branch office environment.

Figure 6 Typical branch office environment

Routing

The switch's routing capabilities include allowing authorized tunneled traffic to securely flow in to and out from the corporation's private network. In addition, the switch can route traffic between two private interfaces, and between its public and private interfaces. The switch can also route traffic from its public interfaces to destinations on the Internet. As a result, you can use the switch to connect your organization to the Internet.

Once you permit traffic to flow between the public and private sides of your network, you also want to ensure that your private network is protected from unauthorized access from the public side. The switch provides a choice of three possible firewall solutions: the Contivity Stateful Firewall, Contivity Tunnel Filter, and Check Point FireWall-1*.

With the addition of an integrated firewall, the switch can perform a variety of secure routing functions, depending upon how you set up the switch's routing capabilities. For example, you can configure the switch to securely route non-tunneled traffic from its private interface, through the firewall, and out its public interface. This configuration would enable users on the private network to access the Internet without requiring a separate, dedicated router.

Both the integrated Contivity Tunnel Filter and the Check Point FireWall-1 provide protection, using features such as packet filtering and antispoofing. The Contivity Tunnel Filter provides:

- VPN routing routes traffic to and from secure tunnels.
- Enhanced routing routes traffic between physical interfaces. This includes traffic between public and private interfaces. Enhanced routing also allows traffic to flow between a tunnel and a public interface.
- Services routing routes traffic used for the services that the switch provides. This type of routing supports tunnel protocols such as IPSec, PPTP, L2TP and L2F. It also supports the use of HTTP and FTP protocols, which are used to manage the switch.

Routing policy service (RPS)

The IP router allows you to control the flow of routing data to and from the routing tables. The routing policy service controls this by providing IP accept and announce policies.

Every IP router maintains a table of current routing information. The routing table manager receives routing updates from the network through the Internet protocols running on the router. Periodically, the routing table manager issues routing updates through the protocols.

Routing table

Like any router, the switch has a routing table that defines how traffic that comes into the switch is routed on to its destination. The routing table can contain both static and dynamic routes. Static routes are manually configured routes that do not change. Dynamic routes, however, do change, as they are learned by using the Routing Internet Protocol (RIP) or Open Shortest Path First (OSPF) from a private interface or a branch office tunnel (the switch does *not* support RIP or OSPF for public interfaces).

About RIP

The Routing Information Protocol (RIP) is a distance-vector routing protocol that allows routers to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring subnets and listen for RIP updates from the routers on those neighboring subnets. Routers use the information in the RIP updates to keep their internal routes current.

About OSPF

OSPF is a link-state routing protocol that maintains a database from which a routing table is constructed from the shortest path, using a minimum of routing protocol traffic. It provides a high functionality open protocol that allows multiple vendor networks to communicate using the TCP/IP protocol family. Some of the benefits of OSPF are:

- Fast convergence
- Variable Length Subnet Masks (VLSM)
- Hierarchical segmentation
- Area routing to provide additional routing protection and a reduction in routing protocol traffic
- Authentication

About VRRP

Virtual Router Redundancy Protocol (VRRP) is one method you can use to configure the switch to maintain a state of High Availability. VRRP is a standard protocol that handles private interface failures. VRRP targets hosts that are configured with static next-hop routing addresses or default gateways. It provides a means of rerouting traffic in the event of a system/interface failure.

Integrated firewall solutions

The switch includes integrated firewall solutions that are designed to meet the needs of a variety of customers. The switch provides a choice of three possible firewall solutions: the Contivity Stateful Firewall*, Contivity Tunnel Filter*, and Check Point FireWall-1*.

With the addition of the Contivity Stateful Firewall, the switch can perform a variety of secure routing functions, depending upon how you set up the switch's routing capabilities. For example, you can configure the switch to securely route non-tunneled traffic from its private interface, through the firewall, and out its public interface. This configuration would enable users on the switch's private network to access the Internet without requiring a separate, dedicated router.

By using stateful inspection, the Contivity Stateful Firewall provides a high level of security, the fastest runtime, and the flexibility to define the rules to fit your environment. The firewall delivers full firewall capabilities assuring the highest level of network security. To do this, the firewall examines packets in both incoming and outgoing directions running against a common security policy. All service rules are interpreted on IP conversations (not packets) and are fully stateful. Security rules do not filter packets directly, but the firewall services determine how to process them based on the security policy defined.

The firewall provides a user interface to help you determine the appropriate rules for your network. The Contivity Stateful Firewall achieves optimum performance as a result of advanced memory management techniques and optimized packet inspection. For further information on the Contivity Stateful Firewall, see *Managing the Contivity Stateful Firewall*.

For many customers, the Contivity Tunnel Filter provides a cost-effective level of protection.


For customers who have Check Point FireWall-1 from Check Point Technologies, LTD, the switch offers an optional, separately licensed integrated firewall. Like the built-in Contivity VPN Switch firewalls, the integrated FireWall-1 is transparent to both users and applications that access the switch.

The relationship between the integrated firewalls and the switch might best be described as a single box that contains the two products in a side-by-side arrangement. All IP packets that go to the switch are passed through the integrated firewall, with the exception of management packets and tunneled packets that are *not* destined for the public network.

Private and public interfaces

The switch provides secure access between your local area network (LAN) and Public Data Networks (PDNs) like the Internet. Private refers to the LAN within your corporation, and Public refers to public data networks. The Public interface accepts only tunneled protocols, while the Private interface accepts both regular (nontunneled) and tunneled protocols. You must be careful to correctly configure each interface of the switch for proper network security.

The LAN interface on the system board is configured to be Private by default. Nortel Networks recommends that its interface be connected to your corporate LAN. This interface accepts regular networking protocols such as TCP/IP, FTP, and HTTP. The Private interface also accepts tunneled protocols (such as IPSec, PPTP, L2TP, and L2F) that can be used for secure management access to the switch. [Figure 7](#) shows private and public interfaces.


Figure 7 Private and public interfaces

Additional interfaces that are inserted into the expansion slots are set to Public by default. Public indicates that this interface is attached to a public data network (PDN) like the Internet. The switch rejects nontunneled protocols and only accepts tunneled protocols like IPSec, PPTP, L2TP, and L2F.

The ability to ping (a program used to send ICMP echo request packets) the Public interface IP address is also supported for diagnostic purposes. Due to security considerations, a public network host is given a 60-packet allotment to establish a connection. If the connection is not successfully established within this allotment, the Public network host must wait approximately 60 seconds before reattempting a tunnel connection. Normally, this allotment is adequate to establish a connection. Beware that continuous pings inhibit your ability to establish a connection.

IP addressing

The following illustration shows sample IP address assignments in a network using a switch. Refer to Figure 8 to see the IP address associations.


Figure 8 Sample IP addressing scheme**Table 2** Sample IP addressing associations

IP address	Description (when applicable, where configured)
192.168.43.6	Dial-up networking to ISP (Internet access, ISP assigned)
192.19.2.30	Public default Internet gateway router
192.19.2.33	Public LAN port IP address (remote user destination address)
192.19.2.32	Firewall public network address
10.2.3.2	Switch management IP address: System→Identity
10.2.3.3	Switch private LAN interface IP address: System→LAN Edit IP address
10.2.3.4	Private network default gateway router: System→Routing Add/Edit Default Route
10.2.3.6	Sample partners FTP server for inventory and price list
10.2.3.7	Firewall private network address
10.2.3.8	DHCP server IP address
10.2.1.1 to 10.2.1.254	Private Network Addresses Assigned to Remote Tunnel Sessions: DHCP pool: Servers→User IP Addr
172.19.2.30	ISP-assigned address

Table 2 Sample IP addressing associations

IP address	Description (when applicable, where configured)
10.2.1.23	DHCP-assigned IP address for a remote user
10.8.4.6	Sample remote user static IP address: Profiles→Users Edit
10.2.4.56	Sample client-specified address: Profiles→Groups Edit IPsec/PPTP/L2TP/L2F

The switch supports the Internetwork Packet Exchange (IPX) protocol. This allows the switch to transmit and receive IPX packets over PPTP, L2TP, and L2F tunnels; not over IPsec, however. The following illustration shows how IPX addresses are associated with private interfaces and a public interface in a typical switch environment. Refer to Figure 9 and Table 3 to see the IPX address associations.

Figure 9 Sample IPX addressing scheme**Table 3** Sample IPX addressing associations

IPX address	Description
0000A100	Public WAN IPX address

Table 3 Sample IPX addressing associations

IPX address	Description
00000B16	802.3 Private LAN, path to NetWare Server
00000C22	SNAP Private LAN, path to Default Nearest Server, an adjacent router, and another NetWare Server
00000D35	802.2 Private LAN, NetWare Server

The switch supports IPX by encapsulating IPX traffic within IP tunnels over PPTP, L2TP, and L2F. The private interfaces and public interfaces can carry IP and IPX traffic simultaneously. The IP addresses are not shown in the preceding illustration.

Tunnels


The switch uses the Internet and remote connectivity to create secure extranets. Remote connectivity through the Public Data Network requires a protocol for safe transport and a connection from the remote user's PC to the PDN. The switch uses the most popular tunneling protocols: IPSec, PPTP, L2TP, and L2F. To form a tunnel, the remote user:

- Establishes a connection with the Public Data Network's point-of-presence (POP), typically through an Internet Service Provider (ISP).
- After the Internet connection is up, the remote user launches a second dial-up connection, and this specifies a connection to a switch. Instead of a telephone number to establish the link, the second connection uses an IP address (or a name if the IP address has been entered into a Domain Name Service server).

This second connection could use either the Point-to-Point Tunneling Protocol (PPTP) or the IP Security (IPSec) tunneling protocol. Tunnels built using L2TP or L2F are slightly different. The tunnel begins at a piece of networking equipment (network access server or NAS) located at the ISP instead of the remote user's PC. The user simply dials into the ISP with a telephone number that causes an L2TP or L2F tunnel to connect directly to a specific corporation. This is similar to a traditional remote-dial service except that the modems are maintained by the ISP and not the corporation.

Figure 10 shows the connection between the remote user and the ISP at the PDN and on to the switch.

Figure 10 Remote user to PDN tunnel and on to the switch



IPSec

The IPSec tunneling protocol is supported by Nortel Networks and other third-party vendors. IPSec is an emerging standard that offers a strong level of encryption (DES and Triple DES), integrity protection (MD5 and SHA), and the IETF-recommended ISAKMP and Oakley Key Determination protocols, and token codes from SecurID* and AXENT*. IPSec offers the following features:

- Client support is available from Nortel Networks and other vendors. No special ISP services are required.
- Support for IP address translation via encapsulation, packet-by-packet authentication.
- Strong encryption and token codes.

Nortel Networks provides the IPSec remote access user client software on the CD that came up with your switch. You can put the client software on a network server for your remote users to download. The client software is a Microsoft* application available for the latest releases of Windows 95*, Windows 98*, Windows* NT, and Windows 2000* Workstation*, and Windows NT* Server. The software comes with complete online Help.

Nortel Networks provides two versions of the IPSec client due to export restrictions. The standard version supports DES (56-bit key) encryption, and the enhanced version supports Triple DES (3DES, 168-bit key). The self-extracting installation files for DES and Triple DES are labeled accordingly on the CD. The installation is simple; the self-extracting installation includes everything necessary to create IPSec tunnels with the switch. For more detail, refer to the readme instructions included as part of the client installation.

PPTP

The Point-to-Point Tunneling Protocol is supported by Nortel Networks, Microsoft, and several other vendors. The Microsoft PPTP client is available for Windows 95, Windows NT Workstation (Version 4.0), and Windows NT Server (except Version 3.51). The Microsoft PPTP client is bundled with Windows 98 software. Network TeleSystems (www.nts.com) provides tunneling product support for Windows 3.1 and Macintosh* operating systems. You can obtain the PPTP client upgrade for Windows 95* directly from Microsoft (www.microsoft.com). Installation instructions are also available from this site.

The PPTP client is on the Contivity VPN Switch CD and is built into the Windows NT operating system. PPTP offers the following features:

- Connections can be made from a range of clients without requiring special ISP services.
- The PPTP client is available for the most common client operating systems.
- PPTP supports IP address translation using encapsulation, support for IPX tunneling, and RC4 encryption (either 56- or 128-bit, within the limits of United States export law).

L2TP

Layer 2 Tunneling Protocol is supported by Nortel Networks, Cisco Systems*, Microsoft, and other vendors. L2TP combines the best features of the L2F and PPTP tunneling protocols. L2TP tunneling allows secure remote access to corporate networks across the public Internet. L2TP tunnels are generally established between a network access server (NAS) at the ISP and the switch. L2TP allows you to specify MS-CHAP, CHAP, or PAP authentication, enable compression, and assign DNS and WINS servers to the tunnel.

You can use support for IPSec transport mode to provide security for L2TP traffic. You can use IPSec transport-protected L2TP tunneling for both remote access traffic and branch office tunnel traffic. Windows 2000 can act as a peer in a branch office connection using L2TP/IPSec or IPSec tunnel mode. Also, Windows 2000 can act as a client using L2TP/IPSec. Authentication for L2TP/IPSEC tunnels can be either shared secret and digital certificate. It also provides configuration support for both voluntary and compulsory L2TP/IPSec remote access connections.

L2TP over IPSEC branch office support

To support the termination of Microsoft Windows 2000 L2TP/IPSec connections for client to gateway connections and gateway to gateway connections, the version 3.01 release has support for IPSec transport mode and uses this to provide security for L2TP traffic.

L2F

The L2F (Layer 2 Forwarding) tunneling protocol is supported by Nortel Networks, Cisco*, Shiva*, and other vendors. L2F tunneling allows remote access to corporate networks across the public Internet. L2F tunnels are generally established between the network access server at the ISP and the switch.

There is no direct client software required for L2F beyond the PPP dialer software, such as the dial-up networking utility provided with Windows 95 and Windows 98. L2F tunnels are actually made from the ISP to the corporate switch on behalf of the user. These connections depend on the domain associated with the


dial-in username. Therefore, ISPs must offer services that are based on L2F; currently, L2F is available on a limited basis. L2F provides IP address translation using encapsulation and support for IPX tunneling, but it does not perform encryption. L2F offers the following features:

- L2F requires special ISP services.
- No requirement for special software on the client.
- No data encryption.

Branch office connections

The branch office feature allows you to configure a secure tunnel connection between two private networks. Typically, one private network is behind a locally configured switch while the other is behind a remote switch. Branch office configuration allows you to configure the accessible subnetworks behind each switch. The configuration also contains the information that is necessary to set up the connection, such as the switch's IP addresses, encryption types, and authentication methods. Local policy restrictions such as access hours, filter sets, and call admission priorities can be applied to limit connectivity into local subnetworks. [Figure 11](#) shows a typical branch office environment.

Figure 11 Typical branch office environment




Network Address Translation (NAT)

NAT is the translation of one network IP address that is used within a LAN to a different IP address that is used outside the LAN. This feature allows a system to be identified by one address on its own network, yet be identified by a totally different address to systems on a different network.

NAT allows branch office connections to eliminate problems with overlapping addresses on both sides of the connection, and it allows you to hide the LAN addresses. [Figure 12](#) shows a sample NAT environment.

Figure 12 Switch sample NAT environment



Authentication

The remote user attempting to dial in to the switch must be authenticated before gaining access to the corporate network. Authentication is one of the most important functions that the switch provides because it identifies users and drives many other aspects of the user-centric functionality.

Because extranet access tends to be used initially with an existing remote access system, the switch is designed to take advantage of the same authentication infrastructure that is used for traditional remote access. An external database can be used for centralized storage of user names and passwords, thereby easing administration. This database can be accessed via an intermediary protocol, such as LDAP or RADIUS.

The switch supports several authentication services and augments them with a group profile mechanism. When a user attempts to get into the network, the switch references a particular group profile to determine encryption strength, filtering profile, quality of service attributes, and more for that user.

This user- and group-specific profile allows users to be grouped together for common attributes, while preserving the flexibility to make exceptions for individual users. The product features and network access that apply to a user can be controlled by the user identity, rather than the source IP address or another mechanism. This is necessary to support mobile users and users coming from other organizations, and to create personal extranets.

Following are details of the authentication services currently supported by the switch, LDAP, or RADIUS.

LDAP

The Lightweight Directory Access Protocol (LDAP) emerged from the X.500 directory service. LDAP is gaining fast acceptance as the directory model for the Internet. Microsoft, Netscape*, and Novell* all support LDAP in their directory service strategies. LDAP is based on directory entries; it has an Internet person schema that defines standard attributes and can be extended to include other attributes. A directory service is a central repository of user information; for example, the switch supports the following elements using LDAP:

- Groups
- Users
- Filters
- Services

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a distributed security system that uses an authentication server to verify dial-up connection attributes and authenticate connections. RADIUS is commonly used for remote access authentication.

Many security systems can be configured with a RADIUS front end to facilitate remote access authentication. RADIUS is also the most common authentication mechanism used by ISPs. Novell NDS*, Microsoft Windows NT Domains, Security Dynamics ACE Server*, and AXENT OmniGuard Defender*, among others, all support RADIUS authentication. Windows NT Domain authentication is used to control access to NT file servers and other resources on NT networks. This is a convenient place for user passwords to be stored, as users likely always remember their file server passwords. Soon, many organizations will use X.509 digital certificates as an authentication mechanism. These certificates work hand-in-hand with public key encryption to identify users to provide a level of assurance that users are who they say they are. Eventually, this type of authentication will be the most common.

SSL and digital certificates

The Secure Socket Layer protocol (SSL) can use Digital Certificates when establishing secure, authenticated connections between SSL clients and servers.

The switch uses a Digital Certificate sent from an SSL-capable LDAP server to authenticate that server. In order for Digital Certificate authentication to succeed, a certificate from the authority certifying the LDAP server must be imported into the switch's certificate store. This type of certificate is often referred to as a Certification Authority (CA) Root certificate.

A single CA Root certificate can be used to certify the authenticity of multiple LDAP servers depending on the organization of your environment's certification hierarchy.

Tunnel certificates

You can use X.509 certificates to authenticate IPSec-based tunnel connections. The switch supports RSA digital signature authentication in the IPSec ISAKMP key management protocol. Remote users can authenticate themselves to the switch using a public key pair and a certificate as credentials. In addition, the switch uses its own key pair and certificate to authenticate the switch to the user.

The switch currently supports the Entrust* product suite (Server: Entrust PKI 4.0 - Web Connector, VPN Connector, and Entrust Authority*; Client: Entrust Entelligence 4.0). Other references pertain to certificate authentication in general.

Using certificates for tunnel connections requires the creation of a *public key infrastructure* (PKI) to issue and manage certificates for remote users and switch servers. The switch software supports the Entrust PKI, Web Connector, VPN Connector, and Entrust client software PKI components. The Entrust software issues certificates for users and the switches.

Encryption

Once you have established a tunnel and have authenticated and authorized users, you might want to secure your data using encryption. Encryption keeps data that is being transmitted private while it crosses a public medium, such as the Internet. Encryption strength varies depending on the encryption method and the bit-length of the encryption keys that are used.

In most cases, encryption methods and the tunneling technologies are linked. For example, PPP within PPTP includes RC4 encryption (either 56- or 128-bit, within the limits of US export law). IPSec can support multiple types of encryption with varying key lengths, such as DES and Triple DES, and token codes from SecurID and AXENT.

The Nortel Networks IPSec tunneling protocol supports the ISAKMP/Oakley protocol that defines how keys for sessions are initiated and updated periodically.

Filters for access control

As you progressively put in place the components for your switch, access control becomes an important security mechanism. You want to have complete control over which users have access to particular servers and services.

Filtering is the mechanism that controls fine-grained access to specific hosts and services. Each user has a specific filter profile based on their group's profile that describes which resources on the network they can access. The filters are defined by:

- Protocol ID
- Direction
- Source, destination IP addresses
- Source, destination port
- TCP connection establishment

A filter profile consists of a list of rules that you create to perform precisely the action that you want. These rules are tested in order until the first match is found. Therefore, the order of the rules is very important. The filtering mechanism works such that if no rule matches then the packet is discarded (denied); this means that no traffic is transmitted or received unless it is specifically permitted.

Management

Management consists of initial and ongoing configuration, status and monitoring, and proactive notification of system events that indicate problems. The switch uses standard HTML Web browsers, Java* scripts, and Java applets.

Since configuration and monitoring tasks are likely to be performed by different individuals, you can assign separate administrators with view only or view and configure access rights.

The switch has a Quick Start configuration option that allows you to set up a default configuration with a single management screen in a few minutes. This allows you to quickly set up a couple of tunnels and verify their operation. The next stage of configuration is the Guided Configuration. This option leads you

through an entire configuration for the switch with section introductions and online Help. Once you become familiar with the navigational menu and submenu options, you can go directly to the Manage Extranet option and browse all of the management screens. Context sensitive online Help is always available by clicking the Help button.

System monitoring provides feedback regarding the status of the system, such as the number of active sessions and the system resources in use.

SNMP traps

The switch supports SNMP MIB II using Gets (status reports). SNMP traps allow you to react to events that need attention or that might lead to problems. You can define through the community name the management stations that receives the SNMP traps. The switch supports all of the SNMP management stations including OPENVIEW*, NetView* 6000, Spectrum*, and Net Manager.

The switch allows the scripting of SNMP alerts so that a combination of system variables can signal an SNMP trap. When a trap occurs, the Nortel Networks icon on a standard management station typically turns red, signaling that there is a problem. The operator would double-click on the icon, which would then open the browser to the Nortel Networks management interface.

Quality of service

With the switch fully configured and many clients dialing in to it, performance and quality of service become important. The switch supports two internal Quality of Service (QoS) mechanisms and can also participate in external network signaling to enhance performance. Forwarding priority allows for prioritized traffic, and Call admission priority allows you to reserve connection resources for high-priority users. In addition, external QoS using Resource ReSerVation Protocol (RSVP) signals the public network to reserve a portion of the network's bandwidth for a specific connection.

Differentiated Services (DiffServ)

DiffServ settings classify and mark packets to receive specified per-hop forwarding behavior on each node along their path. Sophisticated classification, marking, policing, and shaping operations are implemented at network boundaries or hosts. Network resources are allocated to traffic streams by service provisioning policies which govern how traffic is marked and conditioned upon entry to a differentiated services-capable network, and how that traffic is forwarded within that network.

Any DSCPs not recognized are forwarded as if marked for the default behavior, Best Effort (BE).

Bandwidth management

Bandwidth management capabilities let you manage the switch CPU and interface bandwidth resources to ensure that tunneled sessions get predictable and adequate levels of service. Bandwidth management allows you to configure the switch resources for users, branch offices, and interface-routed traffic. Bandwidth components keep track of and control the level of bandwidth being used on the physical interfaces and the tunnels.

Bandwidth management forces tunnels to conform to a set of rates. There are two rates (committed and excess) and excess action (mark or drop). Packets are given different drop preferences depends on whether they are below committed rate (lowest drop preference), between committed and excess rate (higher drop preference), and above excess rate (highest drop preference if excess action is Mark). When there is congestion on the switch, packets are dropped according to their drop preference. When excess action is Drop, all the packets above excess action are dropped.

Forwarding priority

Forwarding priority quality of service allows you to assign each user to one of four priority classes. Each class is guaranteed different maximum forwarding times between the interfaces of the switch. For example, high-priority traffic generated by the company CEO would be protected from high-bandwidth traffic generated by lower-priority users. Or, you might assign the sales team to Priority 1 to make sure that they could always place orders, especially during the quarter-end rush.

The technology that supports Forwarding priority is called *weighted fair queuing with random early detection* (RED). This queuing mechanism gives each of the four user classes (from 1-high to 4-low) a different weight in the amount of service time they receive by the packet-forwarding process. Each class, however, is guaranteed some level of service so that no traffic through the switch is ever completely stalled. It is important to assign users to the four different class levels to make sure that they get the proper service and performance, especially during heavily congested times. QoS is only effective when all associated lines are capable of servicing the forwarding demands at the required speeds.

If a group profile has a Forwarding priority of 1 (highest), it has the highest possible bandwidth guarantee and the lowest level of latency. Packets sent by this group are transmitted immediately even if there is heavy traffic on the switch. Conversely, if a group profile has a Forwarding priority of 4 (lowest), it has the least amount of bandwidth allocated and possibly the highest level of latency. Therefore, fewer packets sent by this group are transmitted while there are higher-level priority packets to be sent when the switch traffic is heavy.

To understand how the Forwarding priority works, the example in Table 4 assumes heavy traffic and a queue of packets. Therefore, packets would be transmitted according to the approximate rates per “pass” that are cited in the table.

Table 4 Bandwidth allocation per priority level

Priority 1	Priority 2	Priority 3	Priority 4
≈60% pass	^a 25% pass	^a 10% pass	^a 5% pass

Of the total packets transmitted in a hypothetical pass, 60 percent would come from the Priority 1 queue; 25 percent from the Priority 2 queue; 10 percent from the Priority 3 queue; 5 percent from the Priority 4 queue.

Call admission priority

Call Admission priority quality of service allows you to assign each user group profile to one of four priority classes (from 1-high to 4-low) for call admission. The switch can reserve connections for each class of user, guaranteeing that a large number of low-priority users do not lock out the high-priority users. When the switch is servicing the maximum number of low-priority sessions, no further low-priority connections are accepted. Once a connection is accepted it is never dropped.

Since there is a maximum number of sessions supported on the switch, it is important to assign users to the proper Call admission priority classes. This ensures that connections are available to the appropriate users when there is heavy utilization. Although other callers may be permitted access to the switch, this access is proportional to the assigned priority level for their group.

By default, any call is admitted access for the first 50 percent of connections regardless of the assigned Call admission priority. The next 25 percent of calls guarantee access to only Priority 1, 2, and 3 callers. The next 15 percent of calls guarantee access to only Priority 1 and 2 callers. For the final 10 percent of calls, only Priority 1 callers are guaranteed access.

For example, assuming a hypothetical maximum of 2000 sessions, Table 5 shows the connections available for each priority based on a percentage of the total capacity:

Table 5 Call admission priority

Capacity	Priority	Available connections
0 to 50%	All	1000
51 to 75%	1, 2, 3	500
76 to 90%	1, 2	300
91 to 100%	1	200

Table 6 shows the maximum number of connections available for each priority:

Table 6 Maximum connections per priority

Priority	Connections
1	2000
2	1800
3	1500
4	1000

RSVP

The switch supports Resource ReSerVation Protocol (RSVP) quality of service for the Internet. Successful external network-level quality of service requires the cooperation of all the devices on the network (between the user and either the access point to the private network or the ultimate destination host). Currently, RSVP is the best-defined technology for resource reservation. However, only a few service providers offer a service that uses RSVP.

The switch is ready to take part in the RSVP signaling that is available in some network backbones and that will increase in the future. The switch signals to the other devices on the public network and describes the level of bandwidth that is needed to ensure adequate performance. This amount of bandwidth is determined by both the data rate that the user has to the Internet, and by the data rate of the link between the Internet and the switch. This beginning stage of RSVP will be extended over time to take advantage of advances in the technology. Meanwhile, you can build RSVP networks using the switch to gain experience in this environment.

The two key components of RSVP are:

- PATH messages, which are constant announcements by the host system or the switch that a certain amount of bandwidth must be kept available.
- RESV messages, which are responses from the client that it wants to reserve the requested bandwidth.

If the client responds to the PATH messages with RESV messages, then RSVP-ready routers attempt the resource reservation. These routers actually reserve the resources requested if they are RSVP-compliant.

Accounting

Accounting is the final component of the switch. Detailed logs record the various activities performed by the switch, including:

- Time and amount of data transferred for each user session
- Security violations and failed authentication
- Device configuration changes, tracking who made the change and when
- System events

These logs are directly available from the management interface, and they can be exported to other applications for additional processing.

The switch supports both internal storage and backup external storage. Automatic backup and archiving of the logs assures you that the logs are available.

As part of your preventive maintenance program, periodically review the various log files for unusual events. Investigate unusual activity and, if necessary, make policy or configuration changes.

Local configuration and LDAP file structures

Configurations available for the switch are stored in two places:

- LDAP database
- Local configuration file

LDAP database information can be shared among multiple switches, while the local configuration file data is specific to a particular switch. For more details, refer to the respective sections of this document.

LDAP database

The LDAP database has been designed to be supported across multiple switches. The LDAP database includes:

- Profiles

- Groups
- Users
- Remote user access
- Filters
- Hours
- Tunnel authentication configuration for IPSec, PPTP, L2TP, L2F
- RADIUS authentication server
- RADIUS accounting
- Remote user IP acquisition, either DHCP or the internal address pool

Configuration file

Local information that is stored in the Nortel Networks configuration file is exclusive to a single switch, including:

- System identity
- Management IP address
- DNS host name
- DNS domain name
- Primary, secondary, and tertiary DNS servers
- Administrator ID and password
- LAN interfaces
 - IP address
 - Subnet mask
 - Default gateway
 - Type
 - Description
- WAN interfaces
 - IP address
 - Description
 - Remote IP address handling, either negotiated or specified
- PPP authentication settings
- PPP advanced settings
- Routing (static routes)
- LDAP
 - Internal, including backup and restoration directory
 - External, including the base domain name, host name or address, port, bind domain name, and bind password for the servers
- User IP address pool
- Administrator tools
- Automatic backup server configurations
- Recovery data
- FTP upgrade access information
- Saved configuration files

- Complete file system
- SNMP traps, including the enabled and disabled traps, host name, community string, trap interval
- System shutdown parameters
- Other
- Alarms

Accounting and logging

Each switch also stores its own accounting and system logging locally. The Event log contains the most recent 2000 events. The System log file is parsed into the Security log and the Configuration log.

Logging and status information are stored in the following:

- Event log
- System log
- Security log
- Configuration log

Command line interface

The command line interface allows you to make configuration changes to the switch via Telnet. You can access the command line interface by initiating a Telnet session to the switch management IP address, for example:

```
telnet 10.0.16.247
```

Then, you must log in using an administrator-privilege user account, for example:

```
Login: admin
Password: *****
%%
```

Upon a successful login, the command line prompt %% appears. You can then enter any of the supported commands.

Serial point-to-point protocol (PPP)

The Serial Point-to-Point Protocol (PPP) feature allows you to manage the switch from a remote location using PPP and the serial interface. If the switch were to become unreachable over the internet, you could still dial up and manage it through the serial interface menu.

With this feature, the serial interface becomes much like a private WAN interface. You can manage through it or even tunnel through it. When configuring Serial PPP, you can set the switch to Auto Detect, or you can specify that either PPP or the Serial menu are the options available through the serial port.

FIPS

You must separately order, purchase, and implement a FIPS kit to be FIPS compliant. This kit contains detailed documentation concerning setting up, operating, and configuring the Contivity VPN Switch 4500 to be FIPS compliant. The FIPS kit also includes tamper-resistant labels to be put on the hardware as instructed in the FIPS kit documentation.

Chapter 2

Managing an Enterprise or Carrier Environment

Managing a large enterprise or carrier network creates management challenges. These environments often consist of equipment from many vendors that support multiple protocols, services, and users. Because network performance is dynamic and the amount of data is large, it can be difficult to detect potential problems in this environment. When problems do arise, it takes time to discover the cause, determine the solution, measure the business impact, and take corrective action. The more time it takes, the greater the negative impact on service availability and overall service quality. While managing service delivery is challenging, the cost of not doing it well can be even greater.

It is important to build network operations centers (NOCs) and operations support systems (OSSs) to effectively manage the distributed, mission-critical networks used to deliver services. The switch allows you to configure a flexible and scalable environment. It also provides you with many tools to help resolve tasks:

- Handling multiple events from several managed resources across several management domains
- Running diagnostic tests
- Isolating problems and their causes
- Determining the impact on users and services
- Notifying management and customer support centers
- Taking corrective actions to resolve each problem
- Removing any failing components from production
- Handling Service Level Agreements (SLAs)

Network configurations

When deploying a large network environment, you probably have one of the following:

- Internet Service Provider (ISP) environment where the ISP manages several different customers
- Enterprise network configuration where a central office manages numerous branch offices

Figure 13 shows a typical ISP environment in which an ISP which provides a managed service to customers and manages several customer deployments from the Network Operations Center (NOC). This allows the ISP to centrally manage both customer A's and customer B's VPNs from the NOC. It assumes all switches in both customers sites have connectivity to the switch NOC through a control tunnel.



Figure 13 ISP network configuration

Figure 14 shows a typical Enterprise branch office environment where the ISP manages its own branch offices from a central office within its enterprise.

Figure 14 Branch office enterprise network configuration

Accessing the Contivity VPN Switch

There are several different ways to access the switch to manage it and obtain information about it. In a large enterprise or carrier environment, you must determine which interface or combination of interfaces best meets your needs.

- Web management interface

You can manage the switch through a Web browser, which connects to the management address of the switch. This allows you to manage the switch from any system without having to install management software. It allows you to configure, access statistical information, obtain reports, and troubleshoot the switch.

- Command Line Interface (CLI)

A command line interface is available from both a serial connection to the switch or from a telnet session. The CLI has a limited set of commands available, which is currently being extended. The goal is to ultimately have a complete CLI from which the user can configure and troubleshoot the switch.

- Bulkload support

In the absence of a fully functional CLI, you can configure certain parameters of the switch through the CLI `load` command. This takes a file with a defined syntax and allows you to easily configure in bulk switch features.

Configuration tasks

Initially, you set up the private interface address, the management address, and the default gateway through the serial interface. This then allows you to manage the device through the Web browser, CLI, or with the bulkload over the private network. Once connectivity is established the initial configuration may be setup.

You can allow access to FTP, DHCP, RADIUS, and DNS servers from the switch through the control tunnel. Control tunnels allow you to easily configure secure tunnels to any switch you want to manage anywhere in the world. This allows you to set up an encrypted tunnel to a customer's switch. Through that tunnel you can perform all the necessary management tasks, such as HTTP, FTP, SNMP, and Telnet.

Enable the appropriate services to be used on the switch:

- Management services -- SNMP, PTP, Telnet, and so on.
- Tunneling protocols -- IPSec, L2TP, PPTP, and so on, with all their parameters, digital certificates, for example.

Configure the servers to support the device:

- LDAP, RADIUS, DHCP, Backup servers, and so on

Configure connectivity and routing:

- Interface addresses, default gateways, and so on
- Static routes or dynamic routing (RIP, OSPF, and so on)
- Packet filters

Configure VPN services:

- Branch office definitions
- Remote access users
- Quality of Service
- Bandwidth management
- User or branch office control tunnel
- SNMP Get and Trap hosts

Configure the systems that will be allowed to manage the switch through SNMP:

- Backup servers for ongoing backups of the system
- Syslog Server to parse syslog messages

These are special tunnels specifically for remote management. With both tunnel types, a secure IPSec tunnel is established to a system that you want to manage. The traffic inside the tunnel is limited to the switch's management IP address only.

The Branch office tunnel allows multiple management systems in the network OPS center to communicate securely with the switch.

A user control tunnels allows a Contivity VPN Client to communicate with a switch that is being managed. This allows network management personnel from anywhere in the world access to the management tasks.

Table 7 shows what tools may be used to configure different parts of the system. In some cases, you may choose to minimally configure the control tunnel on the switch, ship it to the customer site, and then set up the rest of the parameters through the control tunnel.

Table 7 Configuration tasks

Function	Web Interface	Bulkload	Optivity VPN Manager
Configure system:			
DNS server	Y	N	Y
Date/Time/NTP	Y	N	Y
Configure servers:			
LDAP	Y	N	N
RADIUS authentication	Y	N	Y
RADIUS accounting	Y	N	Y
Backup/Restore FTP server	Y	Y	Y
Syslog	Y	Y	N
DHCP server	Y	N	Y
Configure security:			
RADIUS authentication	Y	N	some
Tunneling protocols	Y	N	Y
Certificates	Y	N	N
Firewall/Filters	y	Y-packet	N
Configure management:			
Admin Users	Y	N	N
Management Protocols	Y	N	Y
SNMP Hosts	Y	Y	N
Traps	Y	Y	some
Private/public interfaces	Y	N	N

Table 7 Configuration tasks

Function	Web Interface	Bulkload	Optivity VPN Manager
Configure routing:			
Static routes	Y	N	N
Protocols	Y	N	N
Route policies	Y	N	N
Redundancy	Y	N	N
Network Address Translation	Y	Y	N
Branch office groups/connections	Y	Y	N
Configure user information:			
Groups/Users	Y	Y	N
Hours of access	Y	Y	N
Configure policies:			
Bandwidth management	Y	Y (can apply)	N
DiffServ	Y	Y (can apply)	N

Split tunneling

Split tunneling allows client data to travel either through a tunnel to the enterprise network or directly to the Internet. Although a powerful feature, this could allow an application on the client to maliciously forward packets from the Internet to the enterprise network. Client Policy allows you to determine which network applications and associated protocols and ports a remote user can have active on a workstation while tunneled into the switch. Limiting certain types of network applications from executing while using the split tunneling feature eliminates some security threats. Split Tunneling is disabled by default.

Be careful when choosing the list of network ports that your client can use while using split tunneling because you might prevent acceptable applications from running. When establishing a tunnel, if the client has any network ports open that are not part of the Client Policy list, the tunnel connection is not established and

the remote user is notified. A message is also logged on the switch as to which open port or protocol violated the Client Policy. Network traffic on a client system is monitored constantly to make sure no policy violations occur after the tunnel is established.

Safe mode configuration

The switch can be booted in one of the two system modes: Safe mode or Normal mode. Each mode has its own software image, configuration files, and LDAP database.

A system booted in Safe mode is only allowed to accept secured management tunnel establishment. After the secured management tunnel is established, Telnet, HTTP, and FTP traffic are allowed to come into the switch; no other VPN traffic is allowed through the secured management tunnel or the switch.

In Normal mode, the system operates with the normal software and configuration and transports both VPN traffic and management traffic.

Restricted mode

The Restricted mode feature enforces all management of the switch is done through a secure control tunnel, or the serial menu. Using restricted mode allows you to enforce tighter controls on who can manage the switch. This limits the scope of management to someone who has the proper credentials both to set up the tunnel (if it is an end user) and to login as an administrator (administrative access privileges). Having the proper access privileges in itself acts as a level of security. Additionally, since in restricted mode you are forced to manage the switch through a tunnel, you are guaranteeing data protection through encryption. When the switch is in Restricted Mode, no user on the private side can access the management address of the switch.

Nailed-up control tunnels

Typically branch office tunnels are brought up only when network traffic is destined to go through it. However, if you want the tunnel up all of the time, you can nail it up. You may want to have some control tunnels remain up even when there is no traffic traversing the control tunnel.

Setting up branch office tunnels

When you configure a branch office, you can specify the attributes of the switches that are participating in the connection and set up network parameters for the connection, such as addresses and tunnel type.

When you create a branch office connection, you associate it with a group. The branch office connection then inherits the attributes of that group. You can associate multiple branch office connections with the same group, saving setup time and increasing management efficiency. For example, you might create several VPN connections from various remote sales offices into your enterprise headquarters. You then create all of the connections in the same group so that they all have the same attributes, such as hours of access, encryption method, and password management.

Using bulkload scripts

As a switch administrator, you can automatically generate bulkload scripts. The bulk load command enables you to send a list of commands and parameters to a switch and have them executed in series. You can configure many switches in bulk from a list of settings instead of having to configure each switch manually through the browser interface.

Working with user groups

A group inherits attributes from its parent group. You must explicitly configure a group's unique attributes to override this inheritance. You can assign a group unique network access through packet filtering, attribute support for specific tunneling technologies, minimum encryption levels, authentication mechanisms, access hours, and so forth.

For example, you might want to set up an Administrator group for users who are allowed to manage the switch. This group could be configured to force tunnel connections that use encryption and strong forms of authentication, thereby improving the overall security of the switch. You also might set up a group of Admin_users for users who are allowed to manage users, but cannot manage the switch.

Setting up servers

The switch supports LDAP and RADIUS Authentication Servers. The switch always attempts to authenticate a remote user against the LDAP database. If a User ID and password are found, the switch uses the attributes that are defined for that user's group. The switch can also authenticate against a RADIUS database. When using RADIUS for authentication you can create groups to take advantage of different profiles, or you can simply assign all RADIUS users into a single default group.

Authentication is performed with a protected user ID and password through the ISAKMP key management protocol. Using higher-level encryption, such as Triple DES, decreases performance. The encryption method that you choose determines the level of strength that method provides. All of the encryption methods ensure that the packet came from the original source at the secure end of the tunnel. Some of the encryption types do not appear on non-U.S. models that are restricted by U.S. Domestic export laws. Also, MD5 (Message Digest) provides integrity that detects packet modifications. The following encryption levels are available:

- ESP–Triple DES with MD5 Integrity
- ESP–40- or 56-bit DES with MD5 Integrity
- AH–Authentication Only (HMAC-SHA)
- AH–Authentication Only (HMAC-MD5)

For further information about servers, see [Chapter 6, “Configuring Servers.”](#)

Monitoring the switch

System monitoring provides feedback regarding the status of the system, such as the number of active sessions and the system resources in use.

Configuration file management

Configurations available for the switch are stored in the LDAP database and the local configuration file. The LDAP database information can be shared among multiple switches. The local configuration file contains information specific to a particular switch.

Depending on your environment, you may have to configure not only the system, but the associated servers, security options, management tools, routing, users and policies. The LDAP database provides support across multiple switches.

For further information on configuration files, see [Chapter 9, “Switch Administration.”](#)

System logging

Detailed logs record the various activities performed by the switch, including:

- Time and amount of data transferred for each user session
- Security violations and failed authentication
- Device configuration changes, tracking who made the change and when
- System events

As part of preventive maintenance, you should periodically review the various log files for unusual events. For further information on logging, see [Chapter 9, “Switch Administration.”](#)

SNMP support

You can use SNMP to get statistical information from the switch. You can restrict access control to specific hosts that may retrieve data through SNMP, but by default it is open to all traffic. the switch supports both SNMP V1 and V2c. The list of MIBs supported includes:

- RFC1213 MIB-2
- RFC1724 RIP
- Novell IPX groups
- VRRP MIB (This MIB is still in draft status and has no RFC # yet)

- RFC2178 OSPF
- RFC2667 IP Tunnel MIB

SNMP traps

SNMP Traps alert a management station when certain alarm conditions occur. To receive these traps, the management station IP address must be configured on the switch. Any SNMP management station may be set up to receive the traps; typical systems include HP OpenView and so on. The systems then display an icon representing each switch. Typically this icon may display red if an alarm condition has occurred. When this event occurs, the web browser should be launched against that switch to check the health of the switch, using the Health Check page and checking the event log.

Trap on hardware warnings and alerts

These traps indicate that a warning or alert on any of the following hardware components has occurred. Refer to the discussion on Trap Settings for additional information on specific traps. The default interval for these traps is 00:02:00 (two minutes).

- Intrusion (the top cover has been opened).
- LAN or WAN interfaces.
- One of the dual power supplies has failed.
- Either the critical or normal temperature is out of range.
- One of the voltage indicators is out of range.
- A cooling fan is not working properly.
- System memory is low.
- Disk space is low.

Routing table

The route table contains routes submitted by the routing protocols and the static routes. Dynamic protocols such as OSPF and RIP submit the best route in their view for a specific destination. The switch stores all of the static routes and default routes in the route table. The route table manager chooses the best route based on the following order of protocol priority: direct route, static route, OSPF

route, RIP route, default route. With this and the protocol cost, the route table manager selects the best route and forwards into the forwarding table. This screen provides information about the current routing table, and it allows you to view the routes using filter search criteria and save the information in a file.

You should check the routing tables for any fail-over conditions. If a connection is somehow terminated or lost, the client then attempts to connect to the first listed fail-over switch. It tries each switch in succession and if no connection is established, it stops.

For further information on routing, see [Chapter 5, “Configuring Routing, Firewalls, and IPX.”](#)

Management tasks

After configuring your environment, you can use many of the management tools on the switch to provide you with accounting and statistical information. You should also perform ongoing network management tasks, such as backups and software upgrades and distribution. [Table 8](#) describes ongoing management tasks.

Table 8 Management tasks

Function	Web I/f	Bulkload	VPN Manager
Network Monitoring:			
Alarms/health check	Y	N	N
Accounting:	Y (also RADIUS server)	N	Y
Quality of service configuration	Y	N	N
Ongoing Network Management:			
Backup/restore	Y	Y	Y
Client distribution	Client push	N	N
Software upgrades	Y	N	N

Switch health

The Health Check screen on the switch provides an overall summary of the current state of the switch's hardware and software components. Colored status indicators on the Health Check screen help you evaluate individual component status. Associated hyperlinks allow you to go to screens for corrective action.

The Statistics Memory screen shows how the switch allocates memory, including current free and allocated memory and the cumulative memory. The information includes status, bytes, blocks average and maximum block sizes.

The Status Reports screen allows you to view system and performance data in text or graphical format. You can generate current or historical graphs of valuable system data. The reports feature provides a comprehensive screen or down-loadable reports on user activity. You can view the following report types:

- Administrators--lists users with administrator privileges.
- Users--lists users and the system database groups that they are in.
- System--lists basic system information.
- Sessions--lists session information.
- Failed authorization--lists failed authorization information. You can also generate details for individual listings.
- Expired password--lists users with expired passwords.
- RADIUS diagnostics--lists various RADIUS reports that show whether the switch settings are synchronized with the RADIUS server settings.

Switch data storage

Accounting records are detailed logs that record the various activities performed by the switch. The logs are directly available from the management interface and can be exported to other applications for additional processing.

The Contivity VPN Switch gathers and stores data about the current state of the switch and the connections. The data is stored in files on the switch's hard drive.

- **Session Status: RADIUS Accounting**—the switch stores copies of RADIUS accounting records. These records, which can be retrieved via FTP or sent to a RADIUS server, contain information about each VPN session initiated to the switch.
- **System Data: Data Collection Task**—The data collection task runs on the Contivity VPN Switch and gathers data about the system's status. Each minute, data is captured by the task and written to a data file. The information captured by this task is used to create the graphs and reports available from the Status→Reports page.

RADIUS accounting

The Contivity VPN Switch stores copies of RADIUS accounting records. These records are normally sent to a standard RADIUS Accounting server. To configure a RADIUS accounting server, use the Servers→RADIUS Acct screen.

You can view the information in these records on the Status→Accounting page. These are standard RADIUS accounting records. the switch creates a file for each day and keeps the most recent 60 days of data, which are stored in the SYSTEM/ACCTLOG directory.

Accounting data

A data collection record consists of 16 pairs of entries for each data collection object currently being collected. Each value pair consists of a Field ID and an integer value. The following is an sample data collection record:

0-930057960,1-3,2-3,3-0,4-0,5-0,6-0,7-0,8-0,9-0,10-56,11-76,12-1,13-11021,14-40,15-38,16-0

The following field Ids are currently implemented:

Table 9 Field IDs for data collection records

Field identification	Collected field value	Description
0	TIMESTAMP	Seconds since Jan 1, 1970 - 00:00:00 Hours
1	TOTALSESSIONS	Summary of all sessions
2	ADMINSESSIONS	Number of Admin sessions

Table 9 Field IDs for data collection records

Field identification	Collected field value	Description
3	PPTPSESSIONS	Number of PPTP sessions
4	IPSECSESSIONS	Number of IPSEC sessions
5	L2FSESSIONS	Number of L2F sessions
6	L2TPSESSIONS	Number of L2TP sessions
7	IPADDRESSUSE	Number of IP Addresses in use
8	CPUUSE	Un-filtered CPU usage measurement {integer representing a percent between 0 and 100}
9	CPUSMOOTH	Filtered CPU usage measurement {integer representing a percent between 0 and 100}
10	MEMUSE	Filtered memory usage measurement {integer representing a percent between 0 and 100}
11	BOXPACKETSIN	Number of Inbound Packets
12	BOXPACKETSOUT	Number of Outbound Packets
13	BOXBYTESIN	Number of Inbound bytes
14	BOXBYTESOUT	Number of Outbound bytes
15	BOXDROPPEDPACKETS	Number of discarded packets
16	FAILEDAUTHATTEMPTS	Number of failed authentication attempts
17	LASTFIELDID (this field is never written to data record)	

Accounting records

The data collection system stores records in text-based files stored in the system/ dclog subdirectory. The system stores the most recent 60 days of data. The system stores daily files, summary files, and summary history files.

Daily files contain interval records gathered every 60 seconds. These values are interval values and represent values for the current interval. One record is generated for each 60 second interval and periodically these measurements are flushed to its corresponding daily file. Each day's data is stored in a date-encoded named file, for example 20000622.DC.

Ongoing administration tasks include monitoring the configuration files, backing up and restoring the switch or the LDAP database, and upgrading images and clients.

Service level agreements

Service level agreements (SLAs) are contracts between service providers and customers that define the services provided, the metrics associated with these services, acceptable and unacceptable service levels, liabilities on the part of the service provider and the customer, and actions to be taken in specific circumstances.

The switch supports Quality of Service (QoS) mechanisms that allow you to create SLAs:

- Forwarding priority allows you to prioritize traffic.
- Call admission priority allows you to reserve connection resources for high priority users.
- External Resource ReSerVation Protocol (RSVP) signals the public network to reserve a portion of the network's bandwidth for a specific connection.
- DiffServ allows you to allocate network resources to traffic streams by service provisioning policies, which determines how traffic is marked and conditioned upon entry to a differentiated services-capable network, and how that traffic is forwarded within that network.
- Bandwidth management allows you to manage the switch CPU and interface bandwidth resources to ensure that tunneled sessions get predictable and adequate levels of service. It tracks and controls the level of bandwidth used on the physical interfaces and tunnels.

Backup/restore

You can backup and restore your entire image or your internal LDAP database. The LDAP backup/restore procedure backs up changes to the LDAP LDIF (LDAP Interchange Format) file only. The LDIF file is an intermediate database file that you can use to move data between LDAP servers. Depending on the size of the database, both the backup and restore processes could take extended periods of time.

Upgrades

You can FTP the latest switch software from Nortel Networks. In addition to retrieving software, you can select which version of software to run.

Management applications

In addition to the management tools provided by the switch, you can obtain the VPN Manager product or the Preside suite of products from Nortel Networks.

VPN manager

Optivity VPN Manager* (OVM) is an application designed to simplify and address VPN management challenges. Once installed, OVM is an integrated part of Optivity Network Configuration System* (NCS) 2.0. By using NCS with OVM 1.0, you have a comprehensive, Java-based configuration system and a network management tool for bulk configuration of VPN services. OVM leverages NCS's Web-based, client/server architecture and graphical user interface (GUI) to deliver a system-level approach to efficiently implement and maintain multiple Contivity-based VPN services.

Preside

Preside* is a service-enabling software solution that gives service providers the ability to create and deliver profit-generating services to their customers. It is a network management and operational efficiency tool. The Preside software ties the applications to the network infrastructure and makes it easy to manage on an end-to-end basis. It enables the Internet to become profitable for service providers and gives end-users the personalized services they demand.

Preside offers an open architecture, extensive service-control capabilities, ease-of-use and carrier-grade scalability. It is a fully-functional suite of management applications designed for creation, implementation, and management of high-value, in-demand and revenue-generating applications. It provides the framework to build and retain solid relationships with a customer base looking for competitive, creative and personalized levels of service options.

Preside consists of several service-enabling software applications, which fall into five areas:

- Preside Service Control
- Preside Customer Care and Billing
- Preside Service Activation
- Preside Policy Services
- Preside Service Assurance

The Preside Product Portfolio is also structured to the following solution sets:

- Optical Networks
- Wireless Networks
- Access Networks
- ATM / Frame Relay Networks
- IP Networks and Services
- Circuit Switch Networks
- Application Service Providers

It allows you to create, deliver, and bill for high-value, high-margin services.

You can activate new services, simplify and streamline the delivery of services, and extend carrier management to your existing infrastructure and systems. It can also bring flexibility and growth to your multi-vendor environment by providing the following:

- Synchronize data, optical, telephony, and wireless network capabilities
- Customize network services, applications, and content
- Enable customer self-service over the Internet
- Offer a high quality experience with fast transaction capabilities

Chapter 3

Configuring the Switch

This chapter describes methods that allow you to configure and manage the Contivity VPN Switch.

Full details on hardware installation, including adding local area network (LAN) or wide area network (WAN) cards, are in the *Getting Started* guide that came with your switch. You should complete the hardware installation before starting this chapter.

Table 10 describes the choices you have when first configuring the switch's required parameters. The IP Address Configuration Utility diskette, which comes with your switch, searches for the serial numbers of unconfigured switches. It then displays a table for you to enter the Management IP address, subnet mask, and default gateway (optional).

To configure the switch from the serial interface configuration menu, you must first connect the serial interface cable to the switch. Then you can use a terminal emulation application to enter the Management IP address, subnet mask, and default gateway (optional).

Table 10 Configuration options

Initial configuration method	Result	Advantages and disadvantages
IP Address Configuration Utility (Recommended)	Sets Management IP Address, Subnet Mask, and Default Gateway (optional)	Utility diskette makes initial configuration easy
Serial Interface Configuration Menu (Optional)		Must connect the serial interface cable

Licensing

The only optional software for Version 3.0 is the Advanced Routing feature set, which consists of OSPF, VRRP, Bandwidth Management, and Quality of Service. To enable any of these features, you must first install a license key.

To install a software license key:

- 1** Go to the Admin→Install Keys screen.
- 2** Type the key that you obtained from Nortel Network's Customer Support in the box to the right of Advanced Routing.
- 3** Click on the Install button.

After a valid key is installed, the label "Key Installed" is displayed. It is only necessary to install a key once on each switch.

Click on the Delete button to remove the key. A confirmation message displays and if you click Yes, the key is removed.

IP Address Configuration Utility

Nortel Networks provides a convenient utility to perform the initial configuration of the Contivity VPN Switch. To assign the switch's Management IP address with the IP Address Configuration Utility you must have the following:

- A PC running Windows 95, Windows 98, or Windows NT with a functioning TCP/IP stack. To test the function of your TCP/IP stack, send a PING command to any host.
- The PC must be running on the same subnet as the switch that is to be configured, and it must have an operational network connection.

If your environment does not match these requirements, then you must use the serial interface configuration.

The program ExtNetIP.exe is on a diskette labeled “IP Address Configuration Utility” that accompanies the switch. It is also included on the Contivity VPN Switch CD. You can copy the utility to your hard disk and execute it from there, or you can load the diskette copy from your computer’s diskette drive. The ExtNetIP.exe program launches the IP Address Configuration Utility, which allows you to assign a Management IP address and subnet mask to the switch.

To run ExtNetIP.exe from the diskette:

- 1 Remove the front bezel from 2000 and 4000 series switches.
- 2 Insert the diskette into drive A: and select Start→Run:


a:\ExtNetIP.exe

or, open the My Computer icon on the desktop, open drive A:, then double-click on the extnetip.exe icon:




extnetip.exe

The following screen appears while the program searches for a Nortel Networks switch (and its serial numbers) that has not been configured with an IP address and subnet mask.



- 3 The program automatically enters the Serial Number for the first switch discovered into the table.



- 4 Assign a Management IP address and subnet mask to the switch; the Default Gateway address is optional and can be added later ([refer to “Startup configuration requirements”](#) for descriptions of the required fields).

To add switches to your network, click on Search to automatically add the switch serial numbers. To verify the switches that have been discovered, you can refer to the serial number bar code on the rear of the switch.

- 5 Click Apply to configure the Management IP address, Subnet Mask, and Default Gateway on the switch. The IP address Configuration Utility screen disappears.

After the switch has completed updating its configuration with the Management IP address, Subnet Mask, and optional Default Gateway, your default Web browser automatically opens the switch Welcome screen.



Note: If you move the switch from one network to another, change the Management IP address and Subnet Mask accordingly.

You can obtain help from the extNetIP.hlp file, which is the IP Address Configuration Utility Help file.

Because you need to have Microsoft Internet Explorer*, Version 5.0, the ie5 subdirectory contains the ie5setup.exe file. This file starts the Internet Explorer installation procedure, then links to the Microsoft Web site to download the full Internet Explorer Version 5.0 software. Therefore, you must have Internet access when you run this setup program in order to complete the Internet Explorer installation.

You can install this browser or use a different browser of your choice. Refer to your *Getting Started* for details on getting the Contivity VPN Switch up and running. Nortel Networks recommends that you use the ExtNetIP.exe (the IP Address Configuration Utility). As noted above, you can then assign the necessary IP addressing details to begin managing your switch.

Serial Interface configuration

You can use the following procedure to access the switch from its Serial Interface. The Serial Interface allows you to give the switch a Management IP address and subnet mask so that you can use a Web browser for management.

An alternative to the Serial Interface configuration is to use the Nortel Networks IP Address Configuration Utility, which Nortel Networks recommends for an initial configuration (refer to the Licensing section for additional information).

Your terminal emulator must use the following communications parameters:

- 9600 baud
- 8 data bits
- 1 stop bit
- No parity
- No flow control

The Serial Interface configuration procedure is typically only necessary in a system recovery situation.

- 1 Connect the serial cable (supplied with your switch) from the switch's serial port to a terminal or a communications port of a PC.
- 2 Power on the terminal or PC.
- 3 Using a terminal emulation program on the PC, press [Enter] and you are prompted to supply a user name and password. The factory default user name is **admin** and the password is **setup**.

A menu appears that allows you to enter the following:

- Management IP address
- Management IP subnet mask

- Gateway IP address (optional)
 - Allow HTTP management (check box)
- 4 Follow the screen prompts. For descriptions of the fields required to complete this procedure, refer to “[Startup configuration requirements](#).” A sample screen follows:

Main Menu:

```
1) Interfaces
2) Administrator
3) Private Default Route Gateway*   10.0.0.10
4) Public Default Route Gateway*
5) Create A User Management Tunnel (IPSec)
6) Restricted Management Mode
7) Allow HTTP Management            TRUE
8) Firewall
9) Shutdown
   P) Configure Serial Port
   C) Controlled Crash
   R) Reset System to Factory Defaults
E) Exit, Save and Invoke Changes
```

* Type 0.0.0.0 to delete.

Please select a menu choice (1 - 9,P,C,R,E): 1

- 5 Select E to save your changes and exit.



Note: This administrator’s user ID and password combination is also called the *primary administrator*. This person always has access to all screens and controls, including the serial port and the recovery disk. There can be only one primary administrator.

Safe mode configuration

The switch can be booted in one of the two system modes: Safe mode or Normal mode. Each mode has its own software image, configuration files, and LDAP database.

A system booted in Safe mode is only allowed to accept secured management tunnel establishment. When the secured management tunnel is established, Telnet, HTTP, and FTP traffic is allowed to come into the switch; no other VPN traffic is allowed through the secured management tunnel or the switch.

In Normal Mode, the system operates with the normal software and configuration and transports both VPN traffic and management traffic.

Reconfiguring IP address values

This section describes how to reset the switch IP address values using the serial interface. Use this procedure to:

- Change the Management IP address, Subnet Mask, or Gateway Address to another set of values on the same logical network.
- Move the switch to a different logical or physical network.
- Completely reset the Management IP addresses.
- Reset the switch so that it can be moved to a currently unknown logical network, which allows the network administrator to reuse the Address Configuration Utility to assign new addresses to the switch when it gets there. You might use this option if you do not know the addresses in advance.

Changing IP addresses or moving the switch to another network:

- 1** Log in to the serial interface.
- 2** Select Option 3 (Gateway IP Address). Set the address to 0.0.0.0.
- 3** Select Option 1 (Management IP Address). Set the address to 0.0.0.0.
- 4** Choose E.
- 5** Wait at least 15 seconds. Serial interface data is only written to the hard disk at 15-second intervals.
- 6** Restart the switch. (Use the Reset button on the back for all models except the 1000.)

Changing management IP address when interface IP address is defined

The management IP address must reside on a defined subnet. If there are no defined subnets, the management IP address can reside on any one that you create. You cannot change your management IP address to a subnet different from an already defined subnet.

To change the management IP address when an interface IP address is defined:

- 1 Delete the management IP address (set to 0.0.0.0).
- 2 Set the interface IP address.
- 3 Set the subnet mask.
- 4 Go through the menu again and change the management IP address to the address that you want.

Running the IP address configuration utility

At this point, you have now reset the Management IP addresses and subnet masks. Restart the switch, and you can now use the Nortel Networks IP Address Configuration Utility.



Note: Existing interface values (LAN/WAN), such as users and groups, still reside on the switch. To totally clean up all data in the switch's local database and configuration file, use the Recovery Diskette, reformat the hard drive, restore from the Contivity VPN Switch CD, and then perform the steps listed in the preceding paragraph.

When you are connected to the new network, run the Nortel Networks Extranet Switch IP Address Configuration Utility from a system on the logical network on which you want to install the switch. Assign the new Management IP address and default gateway.



Note: The IP Address Configuration Utility must be run on a system on the same logical network on which the switch resides. The utility must communicate with the switch once the new addresses have been assigned; therefore, it must be running on a system on the same logical subnet.

When you run the IP Address Configuration Utility, consider the following:

- There is a delay of 15 seconds before changes made from the serial interface are written to the configuration file.
- You must change the Gateway IP address before changing the Management IP address (when setting the values to zero to clear them).
- If you run the IP Address Configuration Utility, and it locates a switch, that switch is now waiting for new settings for the Management IP address. It does not respond to future Configuration Utility queries unless you restart the switch.

Startup configuration requirements

Each Contivity VPN Switch is uniquely identified by the system's address and domain name system (DNS) name. The DNS name can be used instead of the IP address to identify the switch and launch its management interface through a Web browser.

The System Identity screen allows you to optionally change your switch Management IP address, and provide the DNS Host Name and Domain Name. Additionally, you can assign up to three DNS addresses to resolve IP address name resolution requests.

You can also reset the switch Management IP address values using the serial interface.

Management IP address

You need the Management IP address to manage all system services from a Web browser, such as HTTP, FTP, and SNMP. This address must be accessible from one of the switch's private physical interfaces. To be accessible, the Management IP Address must map to the same network as one of the private interfaces.

For example, if you plan to assign IP address 10.2.3.3 with subnet mask 255.255.0.0 to the private physical interface, then the Management IP Address must reside in the 10.2.x.x network.

Carefully record the address you enter in the Management IP Address field. Later, during the Quick Start or the Guided Configuration, you are asked to supply IP addresses for the physical interfaces.

Subnet mask

The Subnet Mask field defines how many bits of the IP address represent the network the device is on and how many bits represent the host's ID on the network.

The device uses the value you enter in the Subnet Mask field to determine which IP addresses are directly reachable on the network and which must be routed through a gateway. A sample IP address is 10.2.3.3 with a subnet mask of 255.255.0.0. This indicates that all hosts with addresses 10.2.n.n are directly reachable.

Gateway IP address

The gateway IP address is where packets are routed onto the private network if there is not a specific route in the routing table to the destination.

Web browser

Upon completion of running the IP Address Configuration Utility or using the Serial Interface Configuration, launch a Web browser and enter the Management IP address to invoke the Nortel Networks Login screen. For example, if you assigned the Management IP address as 10.2.3.2, then enter the following Uniform Resource Locator (URL):

```
http://10.2.3.2
```

Select an option in the navigational menu and submenu, and then you are prompted for the Login and Password.

Login and password

Enter the system default Login and Password in lowercase characters, as follows:

Login: **admin**
Password: **setup**

At this point, you should follow the Quick Start Configuration procedure or the Guided Configuration procedure ([refer to “Startup configuration requirements”](#)). Refer to [Web interface configuration options](#) for help in determining which procedure to use.

Preparing for configuration

To properly prepare for installation and configuration of the Contivity VPN Switch, you should have the following items available:

- A Management IP address for the system. You need this address to manage all system services, such as HTTP, FTP, and SNMP.
- An IP address for the LAN port that is available on the system board.
- Any number of Public IP addresses; for example, one IP address for each Public LAN Interface and one IP address for each T1 WAN interface.

- A plan to distribute IP addresses to clients when connections are requested; for example, via a DHCP server or an internal client address pool (with an address pool you need a range of IP addresses).
- An Authentication database: If you are not using internal authentication via the LDAP database then make sure you have either the external LDAP or the RADIUS servers IP address and password or Shared Secret (password).
- An external accounting server, such as RADIUS, with its IP address and Shared Secret (password).
- Client dial-in: Prepare the clients for the type of tunneling protocol they need to use. The PPTP client application is available on the Nortel Networks CD for Windows 95, and it comes with Windows 98 and Windows NT. Nortel Networks also provides the IPsec client on the Nortel Networks CD.
- A complete network topology of the “environment” in which you are testing the switch, including the switch, the default router address, and any other IP addresses that you think might be required.

Table 11 shows the alternatives you have when first configuring your switch. Begin with either the Quick Start or the Guided Configuration. After you are familiar with the switch's Navigational menu and capabilities, you want to select Manage Switch.

Table 11 Web interface configuration options

Configuration type	Results
Quick Start	Configure and test a basic PPTP configuration
Guided Config	Structured switch configuration and management
Manage Switch	Comprehensive switch configuration and management


Welcome screen

The Welcome screen (Figure 15) allows you to enter any of the configuration areas for the Contivity VPN Switch, including:

- Manage Switch
- Manage Notebook
- Quick Start
- Guided Configuration

- Registration

Figure 15 Welcome screen



Before entering the configuration options, you should first register your Switch to activate licenses, warranties, and services.

To start using your switch, choose from one of the following options:

- Click on Manage Switch to begin a configuration management session. This option allows access to all Configuration Management facilities. For your first configuration, follow the Quick Start or Guided Configuration.
- Click on Manage from Notebook to run the Contivity VPN Switch Manager in notebook display mode.
- Click on Quick Start to begin the Quick Start Configuration. This option allows you to configure interfaces, set up PPTP tunnels for up to three users, and establish a connection to the switch. If you prepare for the configuration as recommended, the Quick Start can take as little as 15 minutes to complete.
- Click on Guided Config to begin the Guided Configuration. This option allows access to all Configuration Management facilities. The design and structure of the Guided Configuration, however, is such that you might want to follow the top-to-bottom layout provided. This approach walks you through the entire Navigational Menu from the Profiles to the Admin selections.


Each functional area begins with a summary of the objectives of the area and then steps you through the area (for example, Profiles), one subsection at a time. Context-sensitive help is available at each subsection to supplement the summary.

Provided you have the information required to set up the switch, the Guided Configuration is estimated to take two to three hours to complete, depending on how extensive your configuration is.

- Click on Register Online to register the switch with Nortel Networks. It only takes a few minutes and it gives you access to the latest software and technical tips. Your switch requires Internet access from the private interface in order to register.

The switch's Navigational menu options include the top-level configuration and monitoring areas of the switch. Each of these key areas has secondary levels, which appear once you click on an area; for example, when you click on System, the secondary level listings appear (Figure 16).

Figure 16 Navigational menu



The menu is structured so that you provide system configuration details, followed by profiles for groups and users. You then configure authentication servers, secure tunnels, administrative details, and monitor the status of the switch.


LAN interfaces

The switch provides secure remote access between your LAN and public data networks (PDNs) like the Internet. The network term *private* refers to your corporate LAN and public refers to a public data network. This concept is important because the public interface allows access only through tunneled protocols, while the private interface can support both nontunneled (traditional) and tunneled protocols. Be careful to correctly configure each interface for network security.

The LAN interface that is available on the system board is configured to be private by default. You should connect its interface to your corporate LAN. Additional interfaces that are inserted into the expansion slots are public by default.

The private LAN interface and the Management IP Address must be on the same network, and the public LAN interface should be on a different network, both physically and logically. If your switch has a single network interface and you want to position the switch behind the firewall and router, then you should set the switch's interface type to Private. Figure 17 shows a connection from a LAN to a switch.


Figure 17 LAN-to-switch connection



WAN interfaces

You can assign WAN interface connections between the switch and the PDN. Figure 18 shows the connection attributes that you must configure. These attributes assign WAN interface connections between the switch and the ISP.

Figure 18 Switch-to-PDN configuration



T-1 with integrated CSU/DSU

This section describes the T-1 with an integrated channel service unit/data service unit (CSU/DSU) card for the switch.

A local exchange carrier provides a T-1 service to a customer. The interface between the carrier transmission facility and the customer premises equipment is the CSU/DSU. Nortel Networks now provides a T-1 interface with an integrated CSU/DSU for the switch platform. This is a single peripheral component interconnect (PCI) card with a single T-1 CSU/DSU interface. This card removes the need for you to use an external CSU/DSU.

The T-1 interface with an integrated CSU/DSU conforms to the following standards:

- ANSI T1.403-1995 Telecommunications – Network-to-Customer Installation – DS1 Metallic Interface Specification

- AT&T Technical Reference 62411 ACCUNET T1.5 Service, Description and Interface Specification, December 1990



Note: T-1 cables are required. Although T-1 connectors use RJ-45 8-pins, they are different from 10Base-T connectors. Refer to your getting started guide for additional information.

Configuring the T-1 interface with an integrated CSU/DSU

You can configure your T-1 interface with an integrated CSU/DSU from the System→WAN screen or the serial interface. Following is a list of screens that either allow you to configure or view status for the T-1 interface with an integrated CSU/DSU:

- System→WAN
- System→WAN→Configure
- Admin→Health Check
- Status→Statistics→WAN Status

Newer T-1 services use extended super frame (ESF) framing, which uses out-of-band signaling. The configuration parameters with ESF are:

- Line framing is ESF.
- Line coding is B8ZS.
- HDLC polarity is normal.
- Performance report message value is determined by the T-1 service provider.

Older T-1 services use super frame (SF) framing, which uses in-band signaling. The configuration parameters with SF are:

- Line framing is SF.
- Line coding is AMI.
- HDLC polarity is inverted.
- Performance report message should be set to none as it has no effect in SF framing.

Because SF framing uses in-band signaling, the data can generate a false yellow alarm. These false yellow alarms can be eliminated by setting one fractional T-1 channel to off. If you have the option of using SF or ESF framing, Nortel Networks recommends ESF framing because it provides better diagnostics and does not generate false yellow alarms.

Initial configuration takes place when you install the card, and configuration changes are necessary when adding additional fractional T-1 channels.



Note: You must restart the switch after adding a T-1 card or after enabling a fractional T-1 line.

Verification

You can verify the operational status of your T-1 with an integrated CSU/DSU card by reviewing the Status→Health Check and Status→Statistics→WAN Status screens, and the alarm conditions and LEDs described in Table 12:

Table 12 Alarm conditions and LEDs

	Switch			Telco
State	Receive	LED	Transmit	Receive
Enable	Signal, Framing	Green		
	No Signal, No Framing	Red	Yellow	Alarm
	RAI	Yellow		
	AIS	Blue		
Disable		Off	Blue	Alarm

Alarm conditions

The following alarms are generated when a T-1 with an integrated CSU/DSU card is sending data:

- Yellow – A remote alarm indication (RAI) is generated on the outgoing T-1 signal when the CSU/DSU detects a red alarm, which is a loss of signal or loss of framing.
- Blue – An alarm indication signal (AIS) is generated on the outgoing T-1 signal when the interface is disabled. An AIS is not generated when the interface is enabled.

LEDs

The following LEDs are active when a T-1 with an integrated CSU/DSU card is enabled and is receiving data:

- Green – The interface is functioning normally.
- Red – A red LED indicates a loss of signal or loss of framing by the receiver.
- Yellow – A yellow LED indicates that the card is receiving a remote alarm indication (RAI).
- Blue – A blue LED indicates that the card is receiving an alarm indication signal (AIS).

When all LEDs are off, the interface is disabled.

Troubleshooting

The switch provides the following additional troubleshooting tools, which can be viewed on the Status→Statistics→WAN Status screen:

- Cyclic redundancy check (CRC) errors
- Severely errored frames (SE)
- Frame synchronization bit errors (FE)
- Line code violations (LV)
- Slip events (SL)
- Loop back (LB)

Limitations

The switch does not support the following:

- 7-bit/56K mode, which is referred to as *robbed-bit signaling*. Instead, the switch uses 8-bit/64K, which is referred to as *clear channel*.
- DSX-1 mode
- Channelized T-1

Hardware encryption accelerator

The hardware encryption accelerator introduces a second-generation hardware accelerator card that off loads bulk IPsec encryption duties from the main CPU. This card uses a single HiFn 7811 chip for encryption/compression and is PCI based. It performs triple DES or DES cryptography, LZS compression, and MD5 or SHA-1 authentication. This card is supported for the following switch platforms: Contivity 2500, Contivity 2600, and Contivity 4500. The new hardware accelerator card supports more tunnels encrypted on the card than the first-generation hardware accelerator card.



Note: Version 3.50 is the minimum software revision required for operation and you are limited to one Hardware Encryption Accelerator per Contivity unit.

The hardware encryption accelerator performs bulk encryption and compression algorithms for tunnel traffic in the switch. The encryption accelerator is an add-in device that can perform encryption, compression, or authentication on packets flowing over preestablished tunnels. The accelerator encodes and decodes packets with IPsec ESP encapsulation in hardware using triple DES and DES encryption, LZS compression, and SHA-1 or MD5 authentication. Since the hardware encryption accelerator supports these cryptographic functions, it becomes an element in the overall security envelope provided by the switch. Hardware acceleration is intended to be a background device that requires no additional handling once configured.

The encryption accelerator significantly improves overall capacity of the switch to handle IPsec tunnel throughput. This improvement is realized when 5 or more tunnels are used, and for packet sizes greater than 512 bytes. It enables switch managers to specify (on Profiles→Groups→Connectivity) that users who require high bandwidth are given the highest priority. Then, such users are assured to run

on the accelerator as long as the total number of high priority users does not exceed the maximum sessions of the encryption accelerator card. Also, the switch automatically performs load balancing of tunnel sessions across the accelerator and main CPUs.

The hardware accelerator screen shows the operational status that the switch reports on the hardware accelerator card and allows you to enable automatic recovery in case the card stops running. When the switch detects a recoverable failure, all sessions fail-over and are then handled by the software until the hardware resets and comes back on line.

The accelerator is a user-installable, add-in PCI card for either of the following Contivity VPN switch models:

- 2500
- 2600
- 4500

You must have Administrator privileges to configure the card, which can be installed at any time, and you must restart the switch after configuring it.

The accelerator supports the maximum number of tunnels/sessions for each Contivity platform where it is supported. A tunnel consists of two sessions, one each for incoming and outgoing traffic. Each session comprises a set of logical characteristics and parameters that are associated with a single communication path in a tunnel that renders a full-duplex connection. The number of sessions is double the number of tunnels, since one session is required for each direction and the tunnels are fully bidirectional.

Following is a listing of the switch's configuration, status, and monitoring paths related to the hardware accelerator:

- System → Accelerator → Hardware Accelerator: Configure
- Status → Health Check
- Status → Statistics → Hw Accel Stats
- Status → Statistics → Hw Accel Info
- Status → Event Log

The Hardware Accelerator screen shows the operational status that the switch reports for the card.

Hardware encryption accelerator and load balancing

When the hardware accelerator is installed, the switch's load balancing mechanisms account for accelerator's operation. When assigning tunnels to the hardware accelerator, the switch gives preemptive status to branch office tunnels and high priority user tunnels over non-priority user tunnels.

Hardware encryption accelerator security aspects

As part of the switch's security subsystem, it is imperative that the hardware accelerator does not compromise the overall security envelope of the switch. At startup, whenever the hardware accelerator is manually enabled, or whenever the accelerator recovers from a failure, the Power On self test diagnostics are performed to verify the integrity of the hardware. These tests include validation of the accelerator's encryption, MAC, and compression algorithms against their software counterparts. In the event POST fails, the accelerator is set offline.

Network Time Protocol (NTP)

You can change the time based on your time zone, or make daylight savings time adjustments, as necessary. The switch also provides support for the Network Time Protocol.


The System→Date and Time→Network Time Protocol screen allows you to set up NTP on the switch. NTP synchronizes the clocks of various devices across networks. It also automatically adjusts the time of network devices so that they are synchronized within milliseconds. The switch receives NTP updates from an NTP time server and continuously synchronizes its clock to universal standard time. The switch supports up to eight NTP (unicast) servers and broadcast, multicast servers.

Chapter 4

Configuring tunnels

The switch uses the Internet and tunneling protocols to create secure extranets. The following sections describe configuring the tunnel portion of the switch. The configuration process includes setting up the authentication table and specific tunnel parameters, such as IPSec encryption, L2TP access concentrators and L2F network access servers. [Figure 19](#) shows a typical network illustration with the switch connected to the PDN (public data network) and to a remote user through a tunnel.

Figure 19 Tunnel connection configuration



The connection attributes that you configure in the switch enable the remote user to create a tunnel into the switch. However, you are not configuring the connection from the remote user to the Internet Service Provider (ISP) at this point. The actual connection to the switch is a tunnel that is started from the remote user's PC through its dial-up connection. That connection is to the Internet (typically using an ISP), through the Internet, and ends at the switch on the private, corporate network.

The switch associates all remote users with a group, which dictates the attributes that are assigned to a remote user session. A group can even consist of a single user, thereby creating a personal extranet.

The switch organizes groups in a hierarchical manner. At the top of the hierarchy is the base group. The base group, which might be called "My Company," contains the default characteristics that each new group inherits. You add additional groups to the hierarchy as children of the base group.

The switch does not enforce the Maximum Number of Logins across tunnel types. If you set the number of simultaneous logins to 1, a client can still get another tunnel type connection if the client is configured to use multiple tunnel types. To limit the number of connections a client can have, configure the user for a single tunnel type.

The switch takes precautions against unauthorized users potentially hacking tunneled information when the switch is operating in Split Tunnel mode. The primary precaution is to drop packets that do not have the IP address that is assigned to the tunnel connection as its source address. For example, you establish a PPP dial-up connection to the Internet with an IP address of 192.168.21.3. When you start the tunneled connection to a switch, you are assigned a tunnel IP address of 192.192.192.192. Now, any packets that attempt to pass through the tunnel connection with a source IP address of 192.168.21.3 (or any address other than 192.192.192.192) are dropped. Furthermore, you can enable filters on the switch to limit the protocol types that can pass through a tunneled connection. To completely eliminate security risks, do not use the Split Tunneling feature.




Note: PPP multilink is not supported with branch office tunnels. It is only supported with end user tunnels.

Password aging does not work for administrator accounts. Also, the following are client-specific password management symptoms:

- If you are using the IPSec client, you are warned three times that there will be an impending password expiration. You should change the password immediately. IPSec clients using versions earlier than 1.5.2 do not receive a password expiration warning.
- If you are using the PPTP client with the Extranet Connection Manager, the Extranet Connection Manager generates an impending password expiration.
- Other clients (L2TP and L2F) and PPTP client users who are not using the Extranet Connection Manager have no warning and no longer can log on. You must contact your system administrator if this happens. In this case, the switch is unable to notify the client because it has no actual control over the client. With PPTP, use the Extranet Connection Manager to establish a connection. With L2TP or L2F, set the Password Maximum Age to zero (never expires).

[Figure 20](#) illustrates a group hierarchy. “My Company” is the base group in this example. Research and Development and Finance are child groups of the base group, and they are parent groups to those below them.

Figure 20 Base, parent, and child group associations

Groups are collections of users with the same access attributes and rights. If all users have identical characteristics, then only one group is necessary. You create multiple groups when you need different attributes. A Lightweight Directory Access Protocol (LDAP) database stores users, groups, and their attributes. You can store this database internally (on the switch's hard disk) or externally (on a network host running LDAP server software).

The switch authenticates each user that tries to connect to the switch by checking the user ID and password against a database. The switch supports both LDAP and Remote Access Dial-In User Session (RADIUS) databases for authentication. When using LDAP for authentication, the user is always assigned to a group since LDAP also contains the user, group, and attribute information.

When authenticating a Point-to-Point Tunneling Protocol (PPTP) client against a RADIUS database, the group for a user requesting a session is returned from the RADIUS server as a RADIUS class attribute.

When authenticating an IPSec client, the remote user specifies the Group ID. If the Group ID and Group Password are correct, the switch passes the user ID and Password (or token card) to the RADIUS server for authentication.

You define a set of group attributes and give it a name. This group name is known as the Relative Distinguished Name and it is added to the LDAP database name when performing the database lookup.

In addition to assigning users to groups and providing authentication access, other group characteristics that you can configure include:

- Access hours are time ranges during which access is allowed for users in a group. These time ranges are configured from the Profiles→Hours screen.
- Call admission priority is the priority level (from low to highest) you want to permit for the group. Each level is assigned a percentage of the total number of calls allowed access to the switch. If there is a particularly high number of users logged in, new users could be denied call access, based on their call admission priority, until existing callers disconnect.
- Forwarding priority is the priority level (from low to highest) that you want to provide to sessions for users in this group. Forwarding priority assures a certain level of latency and bandwidth allocation. For example, a group with the highest forwarding priority has the highest possible bandwidth service and the lowest level of latency. Conversely, if there is a particularly high level of traffic on the line, packets for a low priority group might be delayed or dropped. Since a low priority group has the least amount of bandwidth and the highest level of latency, some of its packets would wait until the higher priority level packets have been forwarded or they would be dropped.
- Number of logins is the maximum number of simultaneous logins IPSec clients in the group are allowed.
- Idle timeout is the amount of time a connection can be idle (no data has been transmitted or received through the connection for the specified amount of time). When the idle timeout expires, the session is terminated.
- User attributes include maximum password management, minimum password length, and enabling alphanumeric passwords.
- Tunnel filters are packet filters that control the type of access allowed for users in a group, based on various parameters, including Protocol ID, Direction, IP addresses, Source, Port, and TCP Connection establishment. Choose the Profiles→Filters screen to create tunnel filters.

- RSVP allows you to signal the network for required bandwidth. The client must be configured appropriately for RSVP to work. Only the controlled load-service is supported.
- Tunneling settings allow you to configure IPsec, PPTP, L2TP, L2F settings.

A group inherits attributes from its parent group. For example, if the Research and Development group attributes include All Access Hours and Allow Static Addresses but deny Client-Supplied addresses, PPTP and IPsec tunneling, then the New Products (child) group would inherit these attributes.

Configuring user tunnels

To implement user tunnels, you must configure the following:

- Allowed tunnel access to the switch
- Switch tunneling protocol settings
- A user group
- Adding users to the group

All tunneling protocols are enabled on the public and private networks by default. Since data in tunnels is encrypted, the default setting guarantees that all interactions with the switch are private. To prevent tunnel connections of a particular type (for all users, including Administrators), you can simply disable the tunnel type.

For example, if you want to use IPsec as your only public tunneling protocol, then disable the Public selection for PPTP, L2TP, and L2F. By leaving IPsec, PPTP, L2TP, and L2F enabled on the private side, you can establish tunneled connections to the switch using any of the tunnel types from within your corporation.

To configure tunnel access to the switch:

- 1** Choose Services→Available.
- 2** Select the tunnel type.
- 3** Select the Management Protocol for the switch's private interface.

- 4** Use the RADIUS check boxes to permit RADIUS requests on the public and private interfaces of the switch. If you enable RADIUS traffic, you must also enable RADIUS on the Services→RADIUS screen.

Configuring the switch tunneling protocol settings is dependent on the tunnel type.

- For IPSec, choose Services→IPSec and select the required authentication, encryption and authentication order.
- For PPTP, choose Services→PPTP and select the required authentication and authentication order.
- For L2TP, choose Services→L2TP and select the required authentication, authentication order and configure required L2TP access concentrators.
- For L2F, choose Services→L2F and select the required authentication, authentication order, and configure required network access servers.

To add a user group:

- 1** Go to Profiles→Groups and click on the Add button.
- 2** Enter a Group Name of up to 64 characters (spaces are permitted). For example, you could use Research and Development. The new group is a child of the selected parent group. Therefore, the new group initially inherits the parent group's network access attributes, including authentication, tunnel types, filtering, and priorities. When created, these inherited options can be overwritten for the new group.
- 3** Click Apply and OK to add the group name.

To add a user profile in a group:

- 1** Select a group to which you want to add users from among those in the Group list. If you need to add a new group, select Profiles→Groups.
- 2** After selecting a Group, you must click on Display to view the group members. This allows you to quickly change from viewing one group to another. The last names and first names of the selected group's users appear, sorted by last name.
- 3** Click Add to add a user to the group; the Add User screen appears. This screen allows you to add a User profile. Only options that are enabled for the specified group appear on this screen. Also, only options that the

administrator who is currently viewing the screen has rights to appear. A user profile includes:

- User IDs
- Passwords for the various tunneling protocols
- Assignment of Administrative rights
- An IP address that is always associated with the remote user



Note: You can assign a user to two different groups, but only if the user has two different user IDs. You cannot enter the same user ID in two different groups. When adding a user account, depending on the group configuration, the account can have up to four user IDs. If you are creating an enterprise user ID standard, you should try to avoid schemes that might potentially create conflicts as your company grows. For example, you should not use the user's full first name and last initial.

- 4 Enter the first and last name of the user whose profile you want to add. This is the regular name associated with a person (for example, Mario Smith). This user can have different IDs and passwords for each tunnel type. You can move the user to a another group by selecting a different Group name.
- 5 Enter a Remote User Static IP Address to use in place of a pool (client-specified or DHCP) server-assigned IP address. This IP address is associated with the Static IP Address option in the Groups→Connectivity option (it is only used if the group allows it). If an IP address that is entered here is used instead of a DHCP server-assigned IP address, then only one login is allowed.
- 6 Enter the subnet mask. Assigning the correct subnet mask to a remote IPSec client is important when using Split Tunneling. When you enable Split Tunneling packets destined to a host listed in the Split Tunnel Network list are directed into the tunnel by the IPSec client. All other traffic goes through a standard LAN or dial-up interface. This occurs on the client by adding the routes listed on the Split Tunnel Network list to the route table of the Microsoft TCP/IP stack and pointing those routes to the tunnel adapter interface. A route is also added to the route table based on the subnet mask assigned to the tunnel adapter. The IPSec Subnet Mask field allows you to specifically assign a subnet mask to a remote IPSec client that obtains an IP

address either from the IP Address Pool, DHCP, RADIUS, or a static user configuration.



Note: If a host route for the destination address of the switch exists in the TCP/IP route table prior to launching the Extranet Access Client, the route is deleted when the tunnel is closed.

To search within a selected group and then configure a users account:


- 1 Go to Profiles→Users→UserMangement.
- 1 Select a Group from which you want to search for a particular user from the Group drop-down list box (at the top of the screen), and click Display. The search is limited to the available groups.
- 2 Enter the appropriate text to search for in the input box.
- 3 Select one of the following as the preferred search method, then click Search.
 - Last Name searches for a Last Name. You must enter the entire last name.
 - UID searches for a User ID.
 - Admin Rights searches for anyone who has View or Manage Administrator privileges.
 - LDAP search allows you to enter any LDAP database attribute that is part of the person, organizational Person, or inetOrgPerson object database (for example, cn=common name or sn=surname) to generate the associated user's profile. Refer to your LDAP vendor's documentation for complete details.

Configuring branch offices

The branch office feature allows you to configure a secure tunnel connection between two private networks. Typically, one private network is behind a locally configured switch while the other is behind a remote switch. Branch office configuration allows you to configure the accessible subnetworks behind each switch. The configuration also contains the information that is necessary to set up

the connection, such as the switch's IP addresses, encryption types, and authentication methods. You can apply local policy restrictions, such as access hours, filter sets, and call admission priorities, to limit connectivity into local subnetworks. [Figure 21](#) shows a typical branch office environment.

Figure 21 Typical branch office environment




The section [“Sample branch office configuration procedure”](#) provides sample branch office configurations for two locations, Boston and Cleveland. The initial configurations show connections established with pre-shared keys.

Branch-to-branch with a firewall and a router

In a mixed environment, you might want to tunnel connections to certain networks, and have all other traffic go to the Internet. You must configure the default gateway with a static route to the switch for accessible networks (refer to Profiles→Branch Office→Edit Branch Office Connection). The default private LAN router (the firewall) must redirect packets intended for remote branch office subnets.

In this case, as with any branch-to-branch configuration, you must configure each branch switch with the same encryption settings and pre-shared key (password). Of course, the accessible local and remote subnetwork addresses and subnet masks would be inverted in each switch's configuration.

[Figure 22](#) shows a branch-to-branch configuration with a firewall and a router.

Figure 22 Branch-to-branch with a firewall and a router

In the branch-to-branch illustration, the following interactions take place with a Contivity VPN Switch:

- 1 The PC sends packets to the default route (the firewall).
- 2 The firewall redirects the packets to the local switch branch office connection.
- 3 The encapsulated data goes onto the public LAN.
- 4 The default public LAN route directs the encapsulated data to the remote switch branch office connection.


For switches that have a WAN link, actions 3 and 4 collapse together, and the encapsulated data is directed to the remote server.

Branch-to-branch through a gateway switch

In a three-switch topology, the two indirectly connected switches can create tunnels at will as long as each switch properly includes all of the local and remote subnetworks and subnetwork masks as accessible networks. [Figure 23](#) shows the relationship between three switches and the local and remote networks that must be configured for each link to allow indirectly connected branch offices to bring up tunnels at will. The New York switch in the middle has two branch office connections configured.

All connections must have identical encryption settings. However, only adjacent connections are required to share keys. For example in the following figure, the Boston ↔ New York connection shares keys and the New York ↔ Cleveland connection shares keys. Boston and Cleveland are not required to share keys.

Figure 23 Indirectly connected branch offices



Using NAT


In branch offices, you might have two or more branches that use the same LAN addressing scheme. Nonetheless, users still have to communicate with one another across the branches.

NAT is the translation of one network IP address that is used within a LAN to a different IP address that is used outside the LAN. This feature allows a system to be identified by one address on its own network, yet be identified by a totally different address to systems on a different network.

NAT allows branch office connections to eliminate problems with overlapping addresses on both sides of the connection, and it allows you to hide the LAN addresses (refer to [“Overlapping branch office addresses.”](#))

Figure 24 shows a sample NAT environment.

Figure 24 Sample NAT environment



NAT enables private networks with private addressing to communicate with a public network (Internet) that require public addresses. Typically, companies use private addresses to increase the security of an intranet by hiding the internal IP addresses.

NAT allows privately addressed networks (within an intranet) to use IP addresses that are not assigned to them by the Internet Assigned Numbers Authority (IANA); for example, a 10.n.n.n network address. NAT converts such an internal addressing scheme to an IANA-assigned address before sending a packet out to the public Internet (outside the intranet). This translation generally occurs in a network edge device such as the switch or a router.

In a virtual private network (VPN), NAT allows multiple subnetworks to communicate even though they might be represented by a different address scheme. Using NAT allows you to build an extranet without reconfiguring existing networks. The switch provides NAT for any branch office connection.


Translation types

NAT has certain address translation rules. The two primary types of translation are:

- Static
- Dynamic

A static address is the most specific NAT type or rule and it always overrides dynamic rules. Static addresses are one-for-one. A host name using this rule is always bound to the same external address. For example, 10.2.3.2 within the intranet is always translated to 192.168.34.65. A static address allows bidirectional access to a destination address, as long as you configure the DNS servers accordingly. Additionally, either client device can use the bidirectional address.


[Figure 25](#) shows how source packets from 10.0.0.1 are translated to 192.168.0.1, as statically configured for all packets from the 10.0.x.x network.

Figure 25 Static address translation

Dynamic rules include either port addressing or pooled addressing. The rules are dynamic because an address is assigned to a port or is assigned from a pool, depending on the specific situation and the conditions that currently exist for that situation. Both port and pooled addresses require external end addresses (unlike static addresses, which do not).


Unlike static rules, dynamic port NAT is not one-for-one. All packet transmissions must be initiated from the internal network. For dynamic port translation, the switch checks to see if the packet matches any translation table entries. If an entry exists, then it modifies the destination port and address appropriately. If there are no matching entries, the switch checks to see if the packet is initiating a connection. If so, then it allocates the next available port, adds the address and port to the translation table, and modifies the packet accordingly. It allocates the port assignment from the range of unassigned port numbers. For an incoming packet, if there are no matching entries in the translation table, it drops the packet.

Figure 26 shows how source packets from 10.0.0.1 are translated to 192.168.0.1, through port address translation for all packets from the 10.0.x.x network.

Figure 26 Port address translation

Dynamic pooled NAT is similar to dynamic port NAT. The switch checks to see if an address entry has already been allocated for this situation. If so, it updates the packet addressing and sends the packet. Otherwise, the switch attempts to allocate an address from a pool designated for this session. If an address is available, the switch adds the address pair (the original private address and the newly assigned public address) to the translation table and modifies the packet header. If there are no addresses available, it drops the packet. For an incoming packet, if there are no matching entries in the translation table, then it drops the packet.

Figure 27 shows how source packets from 10.0.0.1 are translated to an address from the pool 192.168.0.1 to 192.168.0.254, through pooled address translation for all packets from the 10.0.x.x network.

Figure 27 Pooled address translation

Overlapping branch office addresses

In branch offices, you might have two or more branches that use the same LAN addressing scheme. Nonetheless, people in the branches still have to communicate with one another.

A typical scenario might include a client on LAN 1 who tries to access the FTP server on LAN 2, and who sends a packet with a source address of 10.0.0.13 and a destination address of 10.0.0.14. Without NAT, the switch would look at the destination address and assume that the destination is on the same LAN as the source device because the addresses are both on the 10.0.0.0 network; therefore, no tunnel connection would be brought up with this packet.

To solve the problem, and allow the client to access the server on the other LAN, you implement NAT on both sides of the branch office connection. In this example, switch 1 defines a remote accessible network of 12.0.0.0, and switch 2 defines a remote accessible network of 11.0.0.0. switch 2 uses a static translation of 10.0.0.14 (server) to 12.0.0.1. switch 1 uses a translation of 10.0.0.13 (client) to 11.0.0.1. As a result, switch 2 must define 11.0.0.0 as the remote accessible network.


With NAT implemented on both sides of the branch office connection, the client can access the FTP server. A packet generated from the client has a source address of 10.0.0.13 and a destination address of 12.0.0.1. switch 1 recognizes that 12.0.0.0 is the remote LAN for the branch office connection. switch 1 translates the source address of the packet to 11.0.0.1 based on the NAT table. switch 2 looks at the destination address of the incoming packet and translate it to 10.0.0.14, but the source address remains 11.0.0.1.

Figure 28 shows the transition of an address from a PC with the IP address of 10.0.0.13 that is translated while being sent to the 10.0.x.x network of a remote branch. Its “source” address becomes 11.0.0.1 while being sent to the destination “12.0.0.0” network. In the figure, the local and remote Accessible Networks for each branch are the inverse of the other branch.



Note: When you have configured two branch offices for NAT, and assigned each office a NAT rule, you must remember that the accessible networks now use the translated addresses.

Figure 28 Overlapping address translation



Creating NAT sets

NAT sets are collections of rules that make up a named set. You can create specific NAT sets for certain conditions, and assign the sets as they are appropriate to the conditions. Typically you apply NAT sets to branch offices that use either static or dynamic address schemes.

The NAT screen allows you to create NAT sets and edit or delete any currently defined NAT sets. To create new NAT sets, define a name and click on Create.

When you edit an existing NAT set, the NAT Rules screen appears. The NAT Rules screen allows you to add new rules or edit existing rules. This screen lists the currently defined rules for a given set.

To create and edit NAT sets:

- 1 Click Profiles→NAT and enter the name of a new NAT set in the create box, then click on Create.
- 2 Click on the newly created NAT name, then click on Edit. The NAT Rules Edit screen appears.
- 3 Click Add Rule. The NAT Add Rule screen appears.
- 4 Select a Translation Type. Next enter the Internal Start Address in the appropriate box. Then enter the External Address and click on OK. The NAT Rules screen appears showing the start and end addresses and the external address.

Sample branch office configuration procedure

This section describes an example of the procedure used to create a branch office VPN connection between two Contivity VPN Switches. In this example, the local switch is at the enterprise headquarters in Boston and the remote switch is at a sales office in Cleveland.

When you set up a branch office connection, you must perform the configuration procedure twice, once for each of the two switches that make up the connection. The branch office settings for the two switches mirror each other. For example, the local address setting that you configure on the Boston switch would be considered the remote address setting when you configure the Cleveland switch.

[Figure 29](#) shows the configuration information and the addresses that are used in this example. It lists the procedure for setting up a branch office connection and the management Web pages that are used during the configuration process. [Figure 30](#) also shows where the information from the figure is entered on the management Web pages.



Figure 29 Sample branch office configuration

Figure 30 Setting up a branch office configuration

As the Administrator of a branch office connection, you can manage the level of access that you give to users of the connection. You specify when the connection is used, what operations can be done through the connection, and which systems on the private networks can be accessed.

Before configuring a branch office connection, check that the following management Web pages are set up in accordance with your management policies and the planned usage for the connection. You use the settings on these pages when you configure of the branch office connection.

- The Profiles→Networks page must list the switch's private networks. In the sample configuration, the local switch's internal network name is `boston_hq` and the subnets are 10.17.20.0 and 10.17.21.0. The remote systems behind the remote switch can reach systems in these networks.

The remote switch's internal network name is `cleveland_sales` and the subnets are 10.17.20.0 and 10.17.21.0. The remote systems behind the local switch can reach systems in these networks.

- The Profiles→Hours page must have the Hours of Access setting that you want to use. The example uses the setting of Anytime.
- The Profiles→Filters page must have the filters that you want to use for the branch office connection. For the example, the local switch uses a filter of permit only dns/http, and the remote switch uses permit all.

Configuring the local switch

You access the Branch Office screen through the Profiles→Branch Office menu selection. Use the subsequent configuration pages to set up a branch office connection. These enable you to specify the attributes of the switches that are participating in the connection and to set up network parameters, such as addresses and tunnel type, for the connection.

To create your new branch office connection, give it a name and associate it with a group. You can choose an existing group or create a new one. The branch office connection then uses that group's attributes, such as password management and encryption. You set the group's attributes on the Profiles→Groups pages.

When you create a branch office connection, you associate it with a group. The branch office connection then inherits the attributes of that group. You can associate multiple branch offices with the same group, thereby saving setup time and increasing management efficiency. For example, you might plan on creating several VPN connections from various remote sales offices into your enterprise headquarters. In this case, you create all of the connections in the same group so they all have the same attributes, such as hours of access, encryption method, and password management.

Use the main Branch Office screen to create new branch office connections and to edit or delete existing connections. You can also add or edit the group that is associated with your branch office connection.

To define a new branch office connection:

- 1** Go to the Define Connection screen to name a new branch office connection and to associate it with a group.
- 2** Go to the Add Group screen to create a new group. You can associate the new group with a branch office connection. You can add a new group that you want associated with the branch office connection. The new group inherits the attributes (for example, Access Hours) of its parent group, which are then used by the branch office connection.
- 3** The list shows all the branch office groups on the switch. Select the group whose attributes you want the new group to inherit. Refer to the Profiles→Groups→Edit→Connectivity screen for additional details on the hierarchical structure of group attributes.
- 4** Enter the Group Name that identifies the new group that you want associated with the branch office connection. The Group Name can be a maximum of 64 characters (spaces are permitted).

Use the initial Define Connection screen to create a new branch office connection. To define a new connection, you first enter its name, then associate the new connection with a group. The connection then inherits the attributes of the group.

To create a new branch office connection:

- 1** Enter the name of the new branch office connection. The name can be a maximum of 64 characters (spaces are permitted).

- 2** The list contains the names of all groups that have been set up on this switch. Select the group that you want to use for the branch office connection. The group is a child of its associated Parent Group and inherits the parent group's network access attributes (refer to the Profiles→Groups→Edit→Connectivity screen for details). You can later modify the new group's inherited options.
- 3** Put a check in the Control Tunnel check box to specify that this branch office connection is for a control tunnel. When you have created and saved a new branch office connection, you cannot later change the control tunnel specification for that connection. For example, if you create a branch office connection that is not a control tunnel, you cannot later change it to be a control tunnel. In this case, you must set up a new connection with the same information as the original connection and create it as a control tunnel.

The Define Connection and the Edit Connection screens allow you to enable the Branch Office feature and to specify routing and networking information, local and remote identification, and authentication attributes for the branch office connection.

To enable a branch office:

- 1** Enter the name you assign to this branch office connection. The name can be a maximum of 64 characters (spaces are permitted).
- 2** Enter the group that defines the attributes that are used by the branch office connection. This group is a child (subset) of its associated parent group and inherits the settings from the parent group. You can click on the Group Details link to view or modify a subset of the group's settings. Modifications of a child group do not change the settings of the parent group.
- 3** Specify the type of routing that you want to be used for traffic going through the branch office connection.
 - If you choose Static routing, you must manually specify the accessible networks (the private internal networks behind a switch that can be accessed via the branch office connection).
 - If you choose RIP, the routing protocol automatically determines the accessible networks based on information that is entered on the System→LAN Interfaces→Edit IP Address screen.
- 4** Click the list to choose the routing type that you want to be used for your branch office connection.

- 5 Click the Enable the Branch Office feature for this switch. As a security mechanism, the Enable Branch Office Connection selection is automatically disabled (the check mark is removed) when you attempt to save an incorrect configuration. For example, if you check the box to enable the branch office connection, then fail to specify the remote address, the Enabled check box is cleared (disabled) and an error message appears when you select the OK button to save your configuration.
- 6 Specify the public interface IP addresses of the switches that form the branch office connection. The Local Endpoint address is the public interface IP address of the switch whose Management Interface you are using. The Remote Endpoint address is the public interface IP address of the switch that forms the opposite end of the branch office connection.
- 7 If you have chosen the Static routing type, this field appears on the screen. It does not appear if you are using RIP routing. The accessible networks are the private internal networks that can be reached through the tunnel connections of this branch office connection.
 - To specify the Local Endpoint networks, click on the list to display available local networks. These networks have been previously set up on the Profiles→Networks screen. The Local networks are the subnetworks on the private internal network of the local switch (the switch whose Management Interface you are currently using).
 - To specify the Remote Endpoint networks, click on Add to choose the Add Networks screen and add the remote networks for the branch office configuration. The Remote networks are the subnetworks on the private network of the remote switch.
- 8 If you choose the Static routing type, this field appears on the screen. It does not appear if you are using RIP routing. NAT allows a system to be identified by one address on its own network, and by a totally different address to systems on a different network. NAT enables you to build your extranet without requiring that you reconfigure or rename your existing network. NAT sets are defined on the Profiles→NAT screen.
- 9 Click the list and select the NAT set that you want to use.
- 10 Select the desired filter that is associated with this connection, or use the default filter of permit all. Packet filtering controls the types of access allowed for users of this branch connection. Filters are based on various parameters, including Protocol ID, Direction, IP addresses, Source, Port, and TCP Connection Establishment. Filters are defined on the Profiles→Filters screen.

- 11 Click the list and choose the filter that this branch office connection that you want to use. The default is permit all. You can specify one filter.
- 12 Use the list to change the tunnel type for the connection. The default type is IPSec. Click the list and select either IPSec, PPTP, or L2TP. If you change the Tunnel Type, the fields in the Authentication portion of this screen change to reflect the different configuration requirements for the new Tunnel Type.

When you configure the authentication that is used between the local and remote branch office switches, the fields that appear depend on whether you are using an IPSec, PPTP, or L2TP tunnel type.



Note: If you create a branch office connection using any IPSec certificate and you choose IP Address as the Alternate name, you must use the IP address of the public interface that is on the branch office end of the connection.

To configure authentication for IPSec:

- 1 Enter the alphanumeric text or hexadecimal string that is used between the local and remote branches for authentication. In order for authentication to occur, you must use the same pre-shared string on both the local and remote branch offices.
- 2 Certificates are associated with each endpoint gateway and allow for mutual authentication between two connections. The certificate portion of the screen includes information about the remote branch office system, the authority that issued the certificate, and the certificate identification. Enter the name of the remote peer initiating the tunnel connection. You can use either a Subject Distinguished Name (Subject DN) or a Subject Alternative Name to uniquely identify the remote branch office system. Specifying both a full subject DN and a subject alternative name on this screen allows the remote peer to use either identity form when making a connection.
- 3 Select a Valid Issuer Certificate Authority from the list. This CA issues the remote peer's certificate or a higher-level CA in the remote peer's certificate hierarchy. The CA must have the trusted flag set on the certificates screen. If a CA hierarchy is being used, all intermediary CAs below the trusted CA must have been imported to the switch. These Certificate Authorities are configured from the System→Certificates Generate→Certificate Request screen.

- 4 If you are using a distinguished name to identify the remote branch office site, you can choose to enter the DN as either a relative distinguished name or a full distinguished name. The DN entered here must exactly match the DN in the remote peer's certificate.
- 5 The Relative distinguished name has the following supported components. Do not include the attribute type as part of your entries in the Relative section. For example, for a name of CN=MyExtranetSwitch, your entry would be MyExtranetSwitch (without the CN attribute type).
 - Common Name with which the server is associated.
 - Organizational Unit (Org Unit) with which the server is associated.
 - Organization with which the server is associated.
 - Locality in which the server resides.
 - State or Province in which the server resides.
 - Country in which the user resides.

You can directly enter the Full Distinguished Name (FDN) in this field rather than entering the individual components in the previously-described Relative distinguished name fields. For example:

CN=MyExtranetSwitch, O=MyCompany, C=US

- 6 You can optionally use a Subject Alternative Name in place of a Subject DN, and specify the format of the name. The following formats are acceptable:
 - Email Name (for example, net_admin@company.com)
 - DNS Name (for example, gateway.cleveland.company.com)
 - IP Address (for example, 192.168.34.21)
- 7 The Local Identity is the name your switch uses to identify itself when initiating or responding to a connection request. You can use either a Subject Distinguished Name (Subject DN) or a Subject Alternative Name to uniquely identify your system. If you select a subject alternative name from your switch's certificate, then that identity is used in place of your switch's subject DN when communicating with peers. Your switch's server certificate only has a subject alternative names if your CA issued the certificate with the alternative names. For example, with the Entrust PKI the VPN connector can issue certificates with DNS names, IP addresses, or Email alternative names.
- 8 Click the list to view all certificates that have been issued to the server. Server Certificates are configured from the System→Certificates Generate→Certificate Request screen.

To configure authentication for PPTP and L2TP:

- 1** Click the list and select the authentication method that you want to use for the branch office connection. When you change the Authentication Type, the screen immediately changes to reflect the requirements of the new authentication method. Any changes that you might have made on the Authentication part of the previous screen are lost.
- 2** Enter the user ID of the local switch that you are configuring.
- 3** Enter the user ID of the remote switch that you are configuring.
- 4** Enter the password for the UID, then confirm the password to verify that you entered it correctly. If you selected a variation of MS-CHAP V2 authentication, no password is required for the Local UID.
- 5** Click to Enable or Disable compression.
- 6** Click to Enable or Disable this selection. This selection is not used if encryption and compression are both disabled.
- 7** The L2TP Access Concentrator field appears if you have selected L2TP as the preferred tunnel type for the branch office connection. Use this entry to specify the L2TP Access Concentrator that performs authentication between the switch and the NAS. When you click on Add in the Accessible Networks section of the Edit Connection screen, the Add Remote Network screen appears.
- 8** Enter the IP address and subnet mask for the new remote network you want to add for the branch office connection.

Configuring the remote switch

After you configure the local switch you must configure the switch located at the remote site. You do this through the Management screens for the remote switch. For this part of the branch office configuration procedure, the Cleveland switch becomes the local switch and the Boston switch is considered remote.

At the Management pages of the remote switch, complete the following steps to define the branch office connection:

- 1** To create a new group, complete these steps:
 - a** Access the Profiles→Branch Office screen.

- b** Click the Add Group button.
 - c** From the list, select the parent group whose attributes the new group inherits, for example /Base.
 - d** Enter the name for the new group. Our sample configuration uses the name cleveland. Click OK to save the settings and return to the Profiles→Branch Office page. You can use the Edit button next to the group name on the Profiles→Branch Office screen to review or modify the group's attributes.
- 2** To define the branch office connection, specify a name for the connection, then associate the connection with a group. Use the following steps:
 - a** At the Profiles→Branch Office screen, click on the Define Branch Office Connection button.
 - b** At the initial Define Connection screen, name the connection. Our example uses the name vpn_to_boston.
 - c** Click the list and associate the connection with a group. The example uses /Base/cleveland. Do not select the Control Tunnel box. Click OK. The configuration page for the new branch office connection displays.

Configuring the remote branch office connection

At the Define Connection screen, enter required information for the branch office connection, for example, Static routing and the IPSec tunnel type.

- 1** Enable the branch office connection by selecting the check box.
- 2** Click the list and choose the routing type that you want to use for your branch office connection. The screen changes to show the appropriate routing fields. This example uses Static routing.
- 3** Specify the addresses of the public interfaces of the two switches forming the connection.
 - a** For the Local Endpoint Address, click on the list and select the address of the local Cleveland switch (132.168.2.3 in our example).
 - b** In the Remote Endpoint Address field, enter the address of the remote Boston switch (132.19.2.30 in the example) that forms the opposite end of the branch office connection.

- 4 Specify the private subnets that can be reached through the tunnel connections of this branch office connection.
 - a For the local network, click on the list and select the address. Our example uses the cleveland_sales entry that you previously created. The local network is the Cleveland switch's private network, which is specified on the Profiles→Networks screen.
 - b For the remote networks, click on the Add button to choose the Add Networks screen. Then add the remote subnetworks' IP Address and subnet Mask. The remote networks are the Boston switch's private networks (10.17.21.0 and 10.17.21.0 in this example). When you are done adding the networks, click on OK to return to the Define Connection configuration page.
- 5 Click on the list and then select the No NAT Translation option.
- 6 To select the filter for the connection, click on the list, then select the filter you want to use (permit all for this example).
- 7 Select the Tunnel Type (IPSec for this example). The screen changes to show IPSec fields.
- 8 Set up the Authentication method for the connection, for example, a Text Pre-Shared Key (refer to IPSec Authentication for additional information). Enter the key of bostoncleveland, then retype it in the Confirm Text String field to verify that you entered it correctly.
- 9 Click the OK button to save the configuration settings for the Cleveland switch.

You have completed the branch office configuration.


Control tunnels

Control tunnels are special tunnels that allow you to securely manage a Contivity VPN Switch over the Internet. The primary reasons for creating control tunnels are secure management and network data integrity. Control tunnels give you secure access to a customer's remote switch so that you can manage it over a network. And control tunnels guarantee that no data from the network behind that customer's switch could ever "leak" through the control tunnel for access by the people on the network that are managing the switch.

You can allow access to FTP, DHCP, RADIUS, and DNS servers from the switch through the control tunnel. Control tunnels allow you to easily configure secure tunnels to any switch you want to manage anywhere in the world. This allows you to set up an encrypted tunnel to a customer's switch. Through that tunnel you can perform all the necessary management tasks, such as HTTP, FTP, SNMP, and Telnet.

Figure 31 shows a sample branch office control tunnel environment where a central VPN server is able to control several VPN devices and configure various services, such as RADIUS, FTP backup, SNMP Traps through Web client management, or Telnet.

Figure 31 Branch office control tunnel




Control tunnel types

There are two types of control tunnels: a branch office control tunnel and a user control tunnel. With both tunnel types, you can establish a secure IPSec tunnel to a system that you want to manage. The traffic inside the tunnels is limited to the switch's management IP address only, which is unique to control tunnels.

Figure 32 shows a special branch office control tunnel from a network operations center in Cleveland and also a user control tunnel.

Figure 32 Sample control tunnel environment



Branch office control tunnels allow anyone on the configured network to communicate with the switch being managed. This allows a switch to communicate with various systems within a company's network operations center or corporate headquarters (the Cleveland private network).

A user control tunnel allows a Contivity VPN Client to communicate with a switch that is being managed. This allows network management personnel from anywhere in the world access to the management tasks.

If you work at a network operations center (NOC) in Cleveland and your job is to manage a customer's switch that is located in Boston, you would want to use control tunnels. On one end of the control tunnel (the switch under management), access is always restricted to the management address only. Access to the Boston switch is limited. The Cleveland end of the tunnel could allow access to its entire private network. This allows multiple systems in your Cleveland NOC to communicate with the management address only of the Boston switch; or for the Boston switch to use remote servers (FTP, DHCP, RADIUS, and DNS servers) on the Cleveland private network.

In this environment, the remote Boston switch a control tunnel to the local Cleveland switch. From any system on the Cleveland network, you can access the management address for the Boston switch. This allows systems on the Cleveland network to initiate management operations on the Boston switch, such as HTTP, FTP, Telnet. Yet because it is a control tunnel, users on the Cleveland private networks cannot exchange packets with users on the private Boston Network.

Additionally, a user control tunnel is configured so that a remote user can establish a control tunnel when using the IPSec client. You create this user account with password authentication in the Control Tunnels group using the serial port.

Restricted mode

The Restricted mode feature prevents management of the switch except through a control tunnel. This limits the scope of management to someone who has the proper credentials both to set up the tunnel (if it is an end user) and to login as an administrator (administrative access privileges). Having the proper access privileges acts as a level of security. Additionally, since in restricted mode you are forced to manage the switch through a tunnel, you are guaranteeing data protection through encryption.

You enable Restricted mode through the Serial Interface menu or the command line interface available through Telnet. In Restricted mode, you can perform the key management functions through the control tunnel, including HTTP, FTP, SNMP, and Telnet. All other attempts to perform these actions outside of the control tunnel will fail.

You cannot enter Restricted mode unless there is an active control tunnel. This ensures there is a mechanism to manage the switch in restricted mode.

Nailed-up control tunnels

You want to have some control tunnels remain up even when there is no traffic traversing the control tunnel. This is generally the case for a branch office versus end user control tunnels. You can do this by creating a script to continuously PING the switch under management. Continuously send PING packets to the management IP address of the switch on the customer premise through the control tunnel from a host at the network operations center. The PINGs must occur at an interval that is less than the Idle Timeout value. These PINGs act as a liveliness detection and perform keepalive signals for the end connection, and report to the sender that the packet was received or that there was no response.



Note: If you change any settings to the branch office connection when using nailed up tunnels, you must bring down the tunnel for the changes to take effect.

Creating control tunnels

To create a special branch office connection, you must create a control tunnel definition on the remote customer switch. There are two methods you can use to create control tunnels:

- On the switch's GUI, use the Profiles→Branch Office screens to create the branch office connection and specify that it is for a control tunnel. Using this portion of the GUI is described in User and Group Configuration.
- For the command line interface, use the switch's command line interface (described in *Reference for the Contivity VPN Switch Command Line Interface*) to set up the connection as a control tunnel. This procedure is described in the following example.

To configure the local switch:

- 1 Initiate a Telnet session to the customer's switch.
- 2 Enter the appropriate control create string, following the required control create parameters already described. A sample string follows:

```
control create boston bostoncleveland 132.19.2.20 132.19.2.30  
192.168.2.3 192.168.20.0 255.255.255.0
```


A special control tunnel filter is used by default with control tunnels to maximize security:

Control Only

3 To view Help, enter control help create.

These are control create parameters that you must enter:

```
CONTROL CREATE <name> <password> <MGMT/Local_P> <Local_endpoint>  
<Remote_endpoint> <Remote_Subnet_Address> <Remote_Subnet_Mask>
```

If you are using the local switch's current Management IP address (132.19.2.20) rather than a substitute, then the network address translation feature is unnecessary. Otherwise, enable control on the remote switch, and enter the control address through the command line interface.

If you enter an address other than the management IP address (MGMT), NAT creates a NAT set with a static rule. The NAT set is called Control plus the name of the connection (for example, Control Boston). This also creates a Network definition that is named Control plus the name of the connection (for example, Control Boston). The network definition contains the NAT management address. In this case, the branch office connection automatically fills in the correct NAT rule and accessible network. When using the Control Create commands, you must enter them in a complete string.

The switch that you are controlling sets a Management only filter by default. This restricts access to the Management IP address only. You can verify your control tunnel connection from the Profiles→Branch Office: Control Tunnels connection field.

After you configure the local switch you must configure the switch located at the remote site. Complete the following steps to define the branch office connection for the remote switch.

To create a new group:

- 1** Access the Profiles→Branch Office screen.
- 2** Click the Add Group button.
- 3** From the list, select the parent group whose attributes the new group inherits; for example, /Base.

- 4 Enter the name for the new group. Our sample configuration uses the name cleveland. Click OK to save the settings and return to the Profiles→Branch Office screen. You can use the Edit button next to the group name on the Profiles→Branch Office screen to review or modify the group's attributes.

To define the branch office connection, specify a name for the connection, then associate the connection with a group:

- 1 On the Profiles→Branch Office screen, click on Define Branch Office Connection.
- 2 On the initial Define Connection screen, name the connection, for example vpn_to_boston.
- 3 Click the name in the list to associate the connection with the group, for example /Base/cleveland.
- 4 Click OK. The configuration screen for the new branch office connection appears.

On the Define Connection configuration screen, you enter required configuration information for the local branch office connection, for example, Static routing and the IPSec tunnel type.

- 1 Enable the Branch Office Connection by selecting the check box.
- 2 Click the list and select the routing type that you want to use for your branch office connection, for example, Static routing. You *cannot* use control tunnels with RIP.
- 3 Click the list and select the routing type that you want to use for your branch office connection, for example, Static routing. You *cannot* use control tunnels with RIP.
- 4 Specify the addresses of the public interfaces of the two switches forming the connection.
 - a For the local endpoint address, click on the list and select the address of the local switch (for example, 132.168.2.3).
 - b In the remote endpoint address field, enter the address of the remote switch (for example, 132.19.2.30) that you want to form the opposite end of the branch office connection.

- 5 Specify the private subnetworks that can be reached through the tunnel connections of this branch office connection.
 - a Click the list and select the network address(es), for example, the cleveland entry previously created. The local network is the switch's private network, which is specified on the Profiles→Networks screen.
 - b Click Add in the Remote Endpoint field, and enter the Management IP address of the local switch and its mask (for example, a host mask of 255.255.255.255).
- 6 Click on the list, then select the No NAT Translation option.
- 7 Click the desired filter for the connection.
- 8 Select the IPSec tunnel type.
- 9 Set up the authentication method for the connection, for example, Text Pre-Shared Key (refer to IPSec Authentication for additional information). Enter the key (for example, bostoncleveland), then retype it in the Confirm Text String field.
- 10 Click OK to save the configuration settings.

Next, you should verify your branch office connection by sending PING packets to the Management IP address of the local switch. Or, you can establish a Web connection to the local switch and attempt to configure it.

Creating a user control tunnel from the serial interface

You can create a user tunnel using the serial interface. Control tunnels allow the management of the switch without access to anything on the network other than the management IP address. This is used to force management through an encrypted tunnel and restricts access to the local resource such as outsourcing management of a switch. You create the Control Tunnel user in the group: /Base/Control Tunnels.

- 1 Connect a serial cable from your computer to the serial port on the switch.
- 2 Connect to the switch using Telnet.
- 3 To create a user account (IPSec), enter the user ID. The user name must be unique within the LDAP database.
- 4 Enter the password and confirm it.

- 5** Enter the IP Address. This is the static IP address assigned to client. The following is the system output:

```
Creating IPSEC[control] account for control
creating Control Tunnels group (this group is /Base/Control
Tunnels with internal ldap)
creating Management Only filter
creating user: control (this username could be anything I just
chose control)
setting IP address 10.1.1.40
creating IPSEC account with uid: control (this password could be
anything I just chose control)
setting password
setting restricted
Success
```

IPSec client features


The switch supports the following IPSec client features:

- Split tunneling
- Third-party IPSec clients
- Forced logout
- Client fail-over
- Client auto connect
- Banner
- Password storage
- Client screen saver
- Domain name
- Client policy

Split tunneling

All IPSec client traffic is tunneled through the switch by default. Split tunneling allows you to configure specific network routes that are downloaded to the client. Only these network routes are then tunneled; any other traffic goes to the local PC interface. Split tunneling allows you to print locally, for example, even while you are tunneled into the switch. [Figure 33](#) shows a sample split tunneling environment.

Figure 33 Sample split tunneling environment



In the previous figure, when you enable Split Tunneling, and you configure split tunnel network IP addresses 10.2.3.4 and 10.10.0.5, when a client establishes an IPSec tunnel, these addresses are loaded into the client application.

The remote user, for example, then downloads his email from the Mail Server at 10.10.0.5, and downloads a document from the Archive at 10.2.3.4. Next, without exiting the tunnel, he can print the document through the PC's local network interface 192.19.2.32 to the Printer at 192.19.2.33. You can enable split tunneling through the Profiles→Groups→IPSec→Edit screen split tunneling field.

You designate which network routes to tunnel through the switch from the Profile→Networks screen. Next, you associate specific network routes to specific groups through the Profiles→Groups→IPSec→Edit screen by configuring the split tunnel networks field.

The switch takes precautions against violators potentially hacking tunneled information when the switch is operating in Split Tunnel mode.

The primary precaution in this release is to drop packets that do not have the IP address that is assigned to the tunnel connection as its source address. For example, if you have a PPP dial-up connection to the Internet with an IP address of 192.168.21.3, and then you set up a tunneled connection to a switch and you are assigned a tunnel IP address of 192.192.192.192, then any packets that attempt to pass through the tunnel connection with a source IP address of 192.168.21.3 (or any address other than 192.192.192.192) are dropped.

Furthermore, you can enable filters on the switch to limit the protocol types that can pass through a tunneled connection.

To completely eliminate security risks, you should not use the Split Tunneling feature.

Third-party IPsec clients

The Client Selection feature enables you to configure your switch to accept tunnel connections from third-party clients, in addition to the Nortel Networks Extranet Access Client. Refer to the *Contivity VPN Switch Release Notes* for a listing of supported third-party clients.

The Client Selection feature provides more flexibility and mobility than was previously available to remote users who want to connect to your switch using a client other than Nortel's Contivity VPN Client. The alternate method of connecting third-party clients requires you to set up a branch office connection, and configure the remote client's IP address as the connection's remote gateway address. This branch office method binds a client machine to a fixed IP address. This can be limiting if a user needs to be able to create tunnels from multiple systems, for example, a work desktop system and a mobile laptop.

With the Client Selection feature, you establish an account for a remote user, rather than for a remote machine. You set up the account within the realm of remote access users, as has always been done for the Extranet Access Client users. This gives the remote user the freedom to create tunnels to your switch from different machines, and from different locations.

When configuring for Contivity VPN Clients, the switch ignores the “Allow undefined networks for non-Contivity clients” field for clients that are not Contivity clients. The switch never allows Contivity VPN Clients to connect to undefined networks. All reachable networks must be defined on the Profiles→Networks screen.

When configuring for clients that are not Contivity VPN Clients, the fields that are preceded by an asterisk are not supported. You must select either the Split Tunneling or “Allow undefined networks field for non-Contivity clients” field for clients that are not Contivity VPN Clients. If you select both, the switch uses the Split Tunneling feature and ignores the “Allow undefined networks” selection.



Note: Nortel Networks recommends that you always specify Split Tunneling for groups used by clients other than Contivity VPN Clients. This ensures that your switch has control over the networks that the third-party client can access.

With Split Tunneling enabled, the third-party clients can only create tunnels to the networks that are listed as split tunnel networks on your switch. If Split Tunneling is disabled and “Allow undefined networks for non-Contivity VPN Clients” is enabled, the clients can connect to all internal networks.

The switch supports both preshared key and RSA digital signature authentication methods. For clients that are not Contivity VPN Clients, you must specify at least one of these authentication methods on the Services→IPSEC screen.



Note: You must ensure that your remote third-party client uses the same Internet Key Exchange (IKE) Phase 1 mode that your switch uses. For Preshared Key authentication, the switch uses IKE Aggressive mode. If the client only supports IKE Main Mode, it must be configured as a branch office due to the IKE restrictions. For RSA Digital Signature authentication, the switch uses IKE Main mode.

RADIUS Authentication is not supported. Also, you can configure a static address for the tunnel from a client other than a Contivity VPN Client, or you can allow the client to use its own IP address as the address used within the tunnel. You can configure the client selection in the Profiles→Groups→Edit→IPSec screen.

Forced logout

For IPSec tunneling, you can specify a time after which all active users are automatically logged off. The default is 0, which means the option is turned off. The possible range is 00:00:01 to 23:59:59.

Client fail-over

Client fail-over uses small packets to check and maintain, or keep alive, the connection between the client and the switch.

Use the Extranet Access Client to disable keep alives between the switch and the client. This option allows you to disable keep alives when tunneling over an ISDN link, since the link is not always active. If an Idle Time-out has been set on the switch, and keep alives have been disabled on the client, the client might not receive notice that the connection has been closed (due to the Idle Time-out), when the physical ISDN connection is not active.



Note: If the idle time-out on the switch logs off the client, and the client has client fail-over configured on the Services→IPSec screen, that client then fails over to the defined failover server, rather than being disconnected as desired.

When the Keep Alive parameter is disabled on the client it prevents the switch and client from exchanging keep alive messages. Therefore, if the connection is lost, the switch does not realize that the client is no longer connected until the idle time is reached. If the Idle Timer can be set to Never, the resulting connection could remain established for a long time, which wastes switch resources.

If the number of Logins is set to 1, which is the default, the client cannot reconnect until the Rekey happens, which by default is in 8 hours. So, if the user has the Disable Keep Alive parameter set on the client, and the connection goes down, the user could be prohibited from reconnecting for 8 hours or more depending on the Rekey value.

Also, do not set the Idle Time-out to 0. If you lose the connection in this situation, you must delete the session from the switch to reconnect.

Client auto connect

The Client Auto Connect feature enables remote extranet clients to connect their IPSec tunnel sessions in a single step. This is similar to the way Microsoft's Dial-Up Networking automatically connects to an ISP when a Web browser is launched. With Auto Connect, extranet client users simply click on the desired destination, for example, a Web page on the private internal network. This first starts their dialup connection, then makes the tunnel connection to the switch, and finally makes the connection to the requested destination. What has, in the past, taken three distinct user operations is now accomplished by a single action.

The Client Auto Connect settings specify those network connections that trigger the Extranet Client's autoconnect feature. For example, you can specify that whenever a remote client attempts to connect to a site in the xyz.com domain, the Client Auto Connect feature is started.

You must make sure that the switch is configured to allow connections to potential destinations. If the switch is not configured properly, the remote user might be able to make the connection to the switch, but cannot access the requested destination. For example, the switch's filters might be set up to deny access to finance.xyz.com, while the Client Auto Connect is configured to start when connections to the xyz.com domain are received. With this configuration, when a remote client tries to access finance.xyz.com, their connection to their ISP and then to the switch is automatically started. However, because of the filters, access to finance.xyz.com is denied.



Note: After you enable the Client Auto Connect feature, you must reboot the PC on which the extranet client is running and manually make sure the extranet client can connect to the switch.

When the Extranet Client successfully connects to the switch, the switch downloads the list of networks and domains that trigger the autoconnect feature. This list, which is stored in the Extranet Client's Registry, is used to determine whether a tunnel connection should automatically be started when one is not already active.

The following client features apply to only the Extranet Access Client:

- Password storage on client

- Client policy
- Client screen saver
- Domain name

Banner

You can customize an enterprise login banner for the Extranet Access Client by entering text into the space provided. This banner appears at the top of the IPSec client upon login.

Password storage

You can allow client systems to save the login password in its password list, or you can require that a remote user enter the password with each request for authentication and access to an IPSec tunnel. Click Enable to allow client systems to save the login password.



Note: When using certificates, saving the password on the client is not allowed.

Client screen saver

Setting this security feature forces the client to use a password in association with the screen saver. When enabled, if the user leaves the system while connected to a tunnel, the system then gets locked out of the tunnel when the screen saver kicks in.

Domain name

This setting enables you to specify the name of the domain used while an IPSec tunnel is connected. Specifying the domain name in this field ensures that domain lookup operations point to the correct domain. This is particularly important for Extranet Clients that use Microsoft Outlook® or Exchange, to ensure that the mail server is mapped to the correct domain.

When a tunnel is connected, the remote client's registry is updated to use the specified domain. When the Extranet Client disconnects the tunnel, the remote client's original domain is again used.

Client policy

Client Policy helps prevent potential security violations that could occur when you are using the split tunneling feature. Split tunneling allows client data to travel either through a tunnel to the enterprise network or directly to the Internet.

Configuring L2TP over IPSec

Windows 2000 supports only L2TP with IPSec transport mode for remote access or branch office. (L2TP cannot be used without IPSec). It supports only RSA Digital Certificates for IPSec transport authentication with the switch. Windows 2000 Professional Server or Advanced Server can act as a Windows 2000 L2TP/IPSec client to a switch server.



Note: After a software upgrade from a version prior to 3.50 to a Version 3.5 or later, IPSec groups created before the upgrade show the new IPSec transport mode connection attribute as enabled with a status of inherited. To correctly set the attribute, you must manually apply the enabled value. The attribute then shows as enabled with a status of configured.

Switch configuration

- 1 Configure an L2TP user account on the switch through the Profiles→Users page and enter an L2TP user ID and password. As with all remote access accounts, you can also enter a static IP address on this page or let the system choose one from an address pool. As an alternative, you can configure a RADIUS server where the L2TP user account resides.
- 2 Before doing any per user configuration, the switch must be issued a certificate and must have the issuer's certificate installed.
 - a Generate a certificate request from the System→Certificates page. This request can be transferred to a CA server that issues the certificate. The certificate can then be installed from the same page.

- b** You must also install the CA server's certificate on the switch through the System→Certificates page. If the Windows 2000 certificate is issued by a different CA, you must also install its certificate.
- 3** Configure an IPSec transport account on the switch in one of three ways:
 - Configure the Subject DN or Alternate Subject Name of the Windows 2000 certificate on the same page as the L2TP user account. Pull down the Valid Issuer Certificate Authority and select the CA who issued the Windows 2000 certificate. Pull down the Server Certificate and choose the switch's certificate to be returned to Windows 2000. Windows 2000 checks the issuer's certificate also, so choose a certificate issued by a CA about which Windows 2000 knows. You may also check the Require Own IPSec Credentials now if you want to ensure this L2TP user always uses this IPSec transport account.
 - Go to the System→Certificates page and select the Enable “Allow All” Feature check box. For the CA that issued the Windows 2000 certificate, select the Allow All Enabled check box. Select a user group from the Default Group pull-down. Be sure the user group selected has Allow IPSec Transport enabled and configured (not inherited) in its IPSec group properties. This configuration is very useful when L2TP user accounts are in RADIUS, since no L2TP or IPSec transport information needs to be stored in the LDAP server per user.
 - Create a separate user that contains the IPSec transport account. Set up the account as described in the first option. Do not check the Require Own IPSec Credentials for either this user or the L2TP user. This configuration is supported mainly for L2TP user accounts that are in RADIUS or compulsory tunneling (many L2TP users sharing an IPSec transport connection). Alternatively (for testing), you could install a single certificate on multiple Windows 2000 PCs, in which case they would be sharing a singled IPSec transport account. Windows 2000 does not support compulsory tunneling.
- 4** Configure the L2TP profile for the user:
 - a** At a minimum, you must set the desired minimum data protection level for the user. L2TP traffic arriving through an IPSec transport that does not meet this requirement is discarded. This is done in the L2TP properties of the group and therefore applies to all L2TP users under this group. No checking is done to determine whether the selection makes sense. For example, selecting 3DES as the minimum protection level implies that 3DES must be able to be negotiated with the Windows 2000 PC. To do

this, DES must be enabled in the Services→IPSec page, must be enabled in the IPSec properties of the group containing the IPSec transport account, and must be configured on the Windows 2000 machine as an acceptable encryption type. Table 13 describes the mapping of minimum data protection levels.

Table 13 Mapping minimum data protection levels to encryption levels

Minimum data protection level	Encryption levels
Triple DES	ESP-Triple DES with SHA1 Integrity-ESP ESP-Triple DES with MD5 Integrity
56-bit DES	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity
40-bit DES	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity ESP-40-bit DES with SHA1 Integrity ESP-40-bit DES with MD5 Integrity
Authentication only	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity ESP-40-bit DES with MD5 Integrity ESP-40-bit DES with SHA1 Integrity ESP-NULL (Authentication Only) with SHA1 Integrity ESP-NULL (Authentication Only) with MD5 Integrity AH-Authentication Only (HMAC-SHA1) AH-Authentication Only (HMAC-MD5)
Not required	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity ESP-40-bit DES with SHA1 Integrity ESP-40-bit DES with MD5 Integrity ESP-NULL (Authentication Only) with SHA1 Integrity ESP-NULL (Authentication Only) with MD5 Integrity AH-Authentication Only (HMAC-SHA1) AH-Authentication Only (HMAC-MD5) Data is allowed through even if it does not come through an IPSec transport with this data protection level.

- b** If the Require Own IPSec Credentials check box is not selected on the L2TP user page, Require IPSec Credentials from Group must select a user group that contains set of allowed IPSec transport accounts. These IPSec transport accounts may be contained at any level below this group. L2TP traffic that arrives through an IPSec transport not contained in this group is discarded.
 - c** Turn on compression in the L2TP group properties if compression is desired. Compression for the PPP traffic is done if both the switch and Windows 2000 agree that compression is enabled. Windows 2000 does not support compression at the IPSec transport level.
 - d** Authentication may be MSCHAPV1, MSCHAPV2, CHAP or PAP. Of these, Windows 2000 prefers to perform MSCHAPV2 followed by MSCHAPV1 followed by CHAP followed by PAP. Windows 2000 does not support L2TP encryption, so if MSCHAP V1 or V2 is enabled, make sure the Not Encrypted check box is also enabled.
- 5** Configure the IPSec transport profile by making sure Allow IPSec Transport is enabled in the group containing the IPSec transport account.
 - 6** By default, Windows 2000 does not have Perfect Forward Secrecy (PFS) enabled. It is enabled by default on the switch. These two settings are not compatible and generate an appropriate error indicating such in the event log when a connection is attempted. To disable PFS on the switch, go to the IPSec properties of the IPSec transport group and disable PFS.

Table 14 User group properties used by IPSec transport connections

Property	Section	Note
Access hours	Connectivity	
Idle timeout	Connectivity	
Client selection	IPSec	Must have clients that are not Contivity VPN Clients enabled
Database authentication	IPSec	Only RSA digital signature is supported with Windows 2000
Encryption	IPSec	
IKE Encryption and Diffie-Hellman Group	IPSec	
Perfect forward secrecy	IPSec	By default, Windows 2000 has this disabled.

Table 14 User group properties used by IPSec transport connections

Property	Section	Note
Forced Logoff	IPSec	
Compression	IPSec	Windows 2000 does not support IPSec encryption. However, this setting does not affect connectivity.
Rekey timeout	IPSec	
Rekey data count	IPSec	
Allow IPSec transport	IPSec	

Windows 2000 configuration

Windows 2000 Professional, Server or Advanced Server may act as a Windows 2000 L2TP/IPSec client to a switch server. The steps for configuring the Windows 2000 side of this follow.

To install a certificate on the Windows 2000 PC, connect to a CA server and get a certificate. For a Windows 2000 Microsoft CA server this involves pointing a browser at the CA server with the following URL:
<IP address>/certsrv.

- 1 Choose Request a Certificate.
- 2 Choose Advanced request.
- 3 Submit a certificate request to this CA using a form.
- 4 On the form provide the Identifying Information. This becomes the subject DN in the certificate that is entered on the switch IPSec transport account.
- 5 Choose IPSec Certificate as the Intended Purpose.
- 6 Select Use local machine store under Key Options.
- 7 When the Certificate has been issued at the CA server, return to the first page
- 8 Choose Check on a pending certificate
- 9 Click Install this certificate. This installs the certificate in the local computer certificate store. To view this store run mmc from the Start→Run prompt. Select Console→Add/Remove Snap-in. From the list of snap-ins choose

Certificates and select Computer account. At the console, expand Personal→Certificates under Certificates (Local Computer). The installed certificate should appear. Clicking on it brings up an information window that indicates its validity and that a private key exists for this certificate.

To install the CA server certificate for the Windows 2000:

- 1 If the switch's certificate was issued by a different CA, that server's certificate should also be installed. For the Microsoft CA, go back to the home page and select Retrieve the CA certificate or certificate revocation list. Click on Install this CA certification path. This installs the CA certificate as a trusted CA, which can be seen in mmc under Trusted Root Certification→Certificates.
- 2 Click on My Computer and click on Network and Dial-up Connections. Click on Make New Connection.
- 3 Choose Connect to a private network through the Internet for the Network Connection Type.
- 4 Enter the interface address of the switch server.
- 5 Edit the properties of this new connection and select the Networking tab. Change the Type of VPN server to L2TP.
- 6 Connect to the switch using the L2TP user ID and password entered on the switch. The certificate installed previously is automatically used to set up the IPSec transport connection.

Branch office

Windows 2000 Server or Advanced Server may act as a Windows 2000 L2TP/IPSec gateway to a switch. Both static routing and dynamic (RIP and OSPF) routing are possible through this branch connection.

To configure the switch:

- 1 Configure an L2TP branch connection on the switch. Go to Profiles→Branch Office.
- 2 Enter the IP Address of the Windows 2000 server as the Remote Endpoint. Select L2TP as the Tunnel Type.
- 3 Choose MS-CHAPV2 unencrypted as the Authentication Type.

- 4** Enter a local UID for the switch.
- 5** Enter a peer UID for Windows 2000.
- 6** Enter a shared password.
- 7** Selected L2TP if you want compression. As with remote access, compression is not supported on Windows 2000 for the IPSec transport connection.
- 8** If you want L2TP tunnel authentication supported, you must provide an L2TP Access Concentrator definition. Windows 2000 does not support L2TP tunnel authentication.
- 9** Select the minimum data protection level. If you select anything other than Not Required, you must set up an IPSec account. Mappings of data protection level to encryption levels are exactly as shown in Table 13.
- 10** As with remote access, the IPSec transport account must be setup. By default, Windows 2000 supports only certificate authentication so a process exactly like that described for remote access must be performed. The CA Allow All authentication option is not available for branch office connections. The L2TP branch office must use the IPSec transport account specified in the connection if data protection is required.
- 11** You can set up routing as either static or dynamic.

To configure Windows 2000:

- 1** You must install a certificate for Windows 2000 and the CA certificates as described in above.
- 2** Start the Routing and Remote Access administrative tool.
- 3** Right click Routing Interfaces and choose New Demand-dial Interface.
- 4** Choose the name of the branch connection. This name becomes the L2TP user ID of the switch. MSCHAPV2 is case sensitive for user IDs. To ensure interoperability with the switch, use lower-case user IDs.
- 5** Select Connect using VPN.
- 6** Select L2TP as the VPN Type.
- 7** Enter the interface address of the CES.
- 8** Select Route IP packets on this interface and select Add a user account so a remote router can dial in.

- 9 Select a password for the switch L2TP user ID. If the switch initiates branch office connections to Windows 2000, this password must match that entered on the switch Branch Office Connection page. If not, then this password does not matter.
- 10 Choose the Windows 2000 L2TP user ID and the shared password. If the Windows 2000 initiates branch office connections to the switch, this password must match that entered on the switch Branch Office Connection page. The Domain field may be left blank.
- 11 The switch supports only MSCHAPV2 as a branch office L2TP authentication method, so be sure you enable this method in the properties (it is by default).
- 12 If you want static routes to demand dial this connection, expand IP Routing→Static Routes and right click New Static Route. Select the interface just created and enter the subnet information. Be sure you enable Use this route to initiate demand-dial connections. Alternatively, you can dial the connection by right-clicking on it and selecting Connect.

Proxy ARP and tunnel to tunnel traffic

You can configure the Contivity VPN Switch to respond to ARP requests on any of the physical interfaces. The switch responds to the following types of routes:

- User tunnels are routes created for user tunnels. This entry is enabled by default and cannot be changed.
- Branch office tunnels are routes available through branch office connections. This option is disabled by default.
- Physical interfaces are routes available through physical interfaces. This option is disabled by default.

To configure Proxy ARP, choose the System→Forwarding screen. You can configure the switch to enable or disable different types of tunnel to tunnel traffic.

Chapter 5


Configuring Routing, Firewalls, and IPX

This chapter describes routing capabilities on the switch and the traffic patterns that are supported. It also describes the options for using firewalls with your switch and Novell's Internetwork Packet Exchange (IPX).

Configuring routing

The switch's routing capabilities include allowing authorized tunneled traffic to securely flow in to and out from the corporation's private network. In addition, the switch can route traffic between two private interfaces, and between its public and private interfaces. The switch can also route traffic from its public interfaces to destinations on the Internet. As a result, you can use the switch to connect your organization to the Internet.

When you permit traffic to flow between the public and private sides of your network, you also want to ensure that your private network is protected from unauthorized access from the public side. The integrated Contivity Stateful Firewall, Contivity Tunnel Filter and the Check Point FireWall-1 provide this protection, using features such as packet filtering and antispoofing. The relationship between the integrated Contivity Stateful Firewall and the switch is shown in [Figure 34](#). The Contivity Tunnel Filter functions in a similar manner.

Figure 34 Traffic flow

Because of the breadth and quality of the routing configurations it supports, in many cases the switch is the only routing device that a corporation needs.

The switch provides the following important routing capabilities. Together, these features make the switch the single solution for the routing needs of most small and medium-sized organizations:

- VPN routing routes traffic to and from secure tunnels.
- Enhanced routing routes traffic between physical interfaces. This includes traffic between public and private interfaces. Enhanced routing also enables traffic to flow between a tunnel and a public interface.
- Services routing routes traffic used for the services that the switch provides. This type of routing supports tunnel protocols such as IPSec, PPTP, L2TP and L2F. It also supports the use of HTTP and FTP protocols, which are used to manage the switch.

Static routes

You can use static routes to set up routes between switches when you do not have any dynamic routing protocol, such as OSPF or RIP. Even if you do have dynamic routing protocols, you may want to use static routes because they provide stronger security. The switch supports multiple default and static routes.

In the absence of any defined route, packets are forwarded to the gateway specified as the default route. These default routes can be either private or public static routes. Private routes are available whether a firewall is enabled. Public routes are available only if an integrated firewall is enabled.

A private default static route is the default route used for traffic that comes into the switch from a private interface. Incoming traffic uses the private default route when there is no public default route defined. If you do not define either a public or private default route, the traffic is dropped. When you add a private default route, the route table adds a new static route.


A public default static route is the default route used for traffic that comes into the switch from a public interface or through a tunnel. If you do not define a public default route, the traffic is dropped. When you add a public default route, a new static route is added to the route table. You can configure multiple default routes to the same destination with different gateways.

You can manually configure static routes on the switch. Based on their states, they are added or removed from the Route Table Manager (RTM). Click Routing→Static Routes to add, edit, or delete static routes.

VPN routing

VPN routing refers to the routing that enables traffic to enter or exit the switch through a tunnel. With VPN routing, the switch supports tunnel-to-private interface traffic and tunnel-to-tunnel traffic, including branch office connections.


Tunneled users might connect to the switch using the Contivity VPN Client or through a branch office connection. This is the traditional VPN traffic pattern that has been supported by the switch since its initial release. If you disable the switch's integrated firewall support, your switch continues to support VPN routing, but does not support enhanced routing or services routing. [Figure 35](#) illustrates the traffic that VPN routing supports.

Figure 35 VPN routing: allowed traffic patterns

Enhanced routing

In addition to the traffic supported by the VPN routing feature, the switch's enhanced routing feature supports traffic patterns between the private network and the Internet, and enables traffic between a tunnel interface and a public interface.

Figure 36 illustrates the variety of routing possibilities supported by the switch. The new traffic type, which is provided by the enhanced routing feature, goes through the integrated firewall. Therefore, one of the switch's firewalls must be enabled to allow the switch to support this type of traffic. The integrated firewall support is an important component in enabling the switch to provide this comprehensive routing environment while maintaining the switch's emphasis on protecting the private network from unauthorized traffic.

Figure 36 Enhanced routing: allowed traffic patterns

Services routing

Services routing describes the traffic patterns for the services that the switch provides. Services routing supports IPSec, PPTP, L2TP, and L2F tunnel protocols. It also helps in management of the switch by supporting HTTP, FTP, and SNMP protocols, and the Check Point FireWall-1 management protocol.

Depending on where it originates, services traffic might enter the switch inside a tunnel or it might be routed using an integrated firewall. Although services traffic can use a tunnel, it differs from the VPN traffic described previously. Services traffic is always destined for the switch, rather than for a system that is on the private internal network behind the switch.


Routing policy service (RPS)

The IP router allows you to control the flow of routing data to and from the routing tables. The routing policy service controls this by providing IP accept and announce policies.

IP accept policies govern the addition of new RIP- or OSPF-derived routes to the routing tables. When RIP or OSPF receives a new routing update, it consults its accept policies to validate the information before entering the update into the routing tables. Accept policies contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matching routes).


IP announce policies govern the propagation of RIP or OSPF routing information. When preparing a routing advertisement, RIP or OSPF consults its area boundary router to determine whether the routes to specific networks are to be advertised and how they are to be propagated. Announce policies contain network numbers (to associate a policy with a specific network) and action information (to specify a route propagation procedure). For OSPF, announce policies are applied only to external routes. For RIP, announce policies apply to all routes, including external routes that are redistributed into RIP and RIP-generated routes.

Every IP router maintains a table of current routing information. The routing table manager receives routing updates from the network through the Internet protocols running on the router. Periodically, the routing table manager issues routing updates through the protocols as shown in [Figure 37](#).

Figure 37 Router configuration with OSPF and RIP

The flow of routing information between the network, the protocols, and the routing table manager is controlled by routing information policies. Each time a routing update arrives from a remote router, the protocol receiving the route consults an accept policy to determine whether to forward the route to the IP routing table manager or drop the route. If the protocol forwards the route to the routing table manager, it adds the route into the routing table as show in [Figure 38](#).

Figure 38 Accept and announce policies



IP0036A

Periodically, the routing table manager announces routes to other routers in the network. The routing table manager forwards a route for advertisement to the protocol. The protocol consults an announce policy to determine whether or not to advertise the route to the network.



Note: OSPF link-state advertisements (LSAs) are received and placed in the link state database (LSDB) of the router. The information in the LSDB is also propagated to other routers in the OSPF routing domain. According to the OSPF standard, all routers in a given area must maintain a similar database. To maintain database integrity across the network, a router must not manipulate received LSAs before propagating them to other routers. To accomplish this, OSPF accept and announce policies act in the following manner:

OSPF accept policies control which OSPF non-self-originated external routing information is passed to the routing table manager. The accept policies control only what the local router uses; they do not affect the propagation of OSPF internal and OSPF non-self-originated external information to other routers.

OSPF announce policies control which self-originated external routing updates are placed into the LSDB for distribution according to the OSPF standard. OSPF announce policies affect what other routers learn but only with regard to the local router's self-originated information.

To configure routing policy services, go to Routing→Policy.

Routing table

Like any router, the switch has a routing table that defines how traffic that comes into the switch is routed on to its destination. The routing table can contain both static and dynamic routes. Static routes are manually configured routes that do not change. Dynamic routes, however, do change, as they are learned by using the Routing Internet Protocol (RIP) or Open Shortest Path First (OSPF) from a private interface or a branch office tunnel (the switch does *not* support RIP or OSPF for public interfaces).

The route table contains routes submitted by the routing protocols and the static routes. Dynamic protocols such as OSPF and RIP submit the best route in their view for a specific destination. The switch stores all of the static routes and default routes in the route table. The route table manager chooses the best route based on the following order of protocol priority: direct route, static route, OSPF route, RIP route, default route. With this and the protocol cost, the route table manager selects the best route and forwards into the forwarding table.

The switch's routing table can be thought of as having two separate parts. One part contains the routes for traffic that uses the switch's public interfaces and a second part of the table has the routes for traffic using the private interfaces.


The following list shows the types of routes that might be in the switch's routing table:

- Static routes
 - To public interfaces
 - To private interfaces
 - To a configured branch office
- Dynamic routes
 - To private interfaces
 - To a configured branch office
- Default routes
 - To public interfaces
 - To private interfaces
- Host routes
 - Routes added for dialed VPN users (for example, Contivity VPN Clients or PPTP clients)
- Utunnel Routes
 - Host/network routes for clients that log in using the client address redistribution feature

The following figure further illustrates how the types of routes might be contained in the switch's routing table as well as the traffic that uses each part of the route table. Public routes are routes that are available only if an integrated firewall is enabled, while private routes are available whether a firewall is enabled.

To view the Route Table, go to the Routing→Route Table screen. [Figure 39](#) shows the switch's routing table public and private traffic.

Figure 39 Switch routing table



Route lookup

When a packet arrives, the switch performs a full lookup in its routing table. The following sequence is used to determine which route to use:

- If the switch's firewall support is enabled, all the public and private routes in the routing table are available to the traffic.
- If firewall support is not enabled, then only the private portion of the routing table is available.
- If the traffic's destination route is not found in the table, the table's public and private default routes are invoked as described in the following section.

Default routes

Default routes are used when the switch receives traffic for which no matching route exists in the routing table. The use of the default routes depends on several factors, such as whether integrated firewall support is enabled and where the traffic originated (for example, from the public or private interface). The table “[Traffic patterns](#)” lists how the use of default routes is determined.

The default gateways are the addresses of the next hop router. Packets are routed through the default gateway onto the private or public network when the Routing Table does not have a specific route to the destination.

Rules of redistribution for RIP and OSPF

Table 15 describes the rules of redistribution for RIP and OSPF with the firewall enabled or disabled.

Table 15 Redistribution rules

Redistributed Route	Firewall ON	Firewall OFF
Public Direct Route	Yes	No
Public Static Route	Yes	No
Public UTunnel Route	Yes	No
Private Direct Route	Yes	Out physical - No; out tunnel - Yes
Private Static Route	Yes	Out physical - No; out tunnel - Yes
Private UTunnel Route	Yes	Out physical - No; out tunnel - Yes
Tunnel Static Route	Yes	Out physical - Yes/No Out other tunnel - Yes/No Yes in general, but RIP has more specific control per interface to determine whether the static routes are redistributed.
Tunnel Dynamic Route	Yes	Out physical - Yes; out other tunnel - Yes

RIP redistributes static and default routes by default if Route Policy is Enabled in Version 3.0 or higher of the switch. To avoid any problems that this could cause, you can add a route policy on any (or both) switches to deny the announcement of Direct routes using the Routing Policy screen.

About RIP

The Routing Information Protocol (RIP) is a distance-vector routing protocol that enables routers to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring subnets and listen for RIP updates from the routers on those neighboring subnets. Routers use the information in the RIP updates to keep their internal routes current.

For RIP, the best path to a destination is the path with the fewest hops. RIP computes distance as a metric, usually the number of hops (or routers) from the origin subnet to the target subnet. RIP can handle a maximum of 15 hops.

RIP is one of the most common interior gateway protocols used in the Internet. RIP Version 2 is upward compatible with RIP Version 1 and corrected many of the RIP Version 1 shortcomings, such as subnet routing, authentication, support for multiple autonomous systems, and multicast support for route messages.

The switch supports RIP for routing traffic within the private network and between branch office connections. RIP is used only on the switch's LAN interfaces, because the switch's WAN interfaces are public interfaces. The switch supports both Version 1 and Version 2 of the RIP protocol. For additional information on RIP, refer to the RFCs, which are on the Internet Engineering Task Force (IETF) Web site at www.ietf.org.

To configure RIP global parameters, go to Routing→RIP. When globally enabled, RIP can be added to specific interfaces by going to the Routing→Interfaces screen.

- RFC 1058 – Routing Information Protocol: Describes the Routing Information Protocol (RIP), which is loosely based on the program “routed,” distributed with the 4.3 Berkeley Software Distribution. The specifications in this RFC represent a combination of features taken from various implementations of this program.
- RFC 1721 – RIP Version 2 Protocol Analysis: Describes the key features of the RIP Version 2 protocol and the current implementation experience.

- RFC 1722 – RIP Version 2 Protocol Applicability Statement: Describes how RIP Version 2, which is an extension to RIP Version 1, may be useful within the Internet.
- RFC 1723 – RIP Version 2 Carrying Additional Information: Specifies an extension of the Routing Information Protocol (RIP) that expands the amount of useful information carried in RIP messages and that adds a measure of security.

Using RIP on your switch

The switch sends RIP broadcast/multicast messages at regular intervals. These messages contain information about routes that the switch can reach. Other routers on the network pick up these messages, and in turn update their routing tables and then send out route messages to their peer routers. The switch's RIP support allows you to:

- Enable/disable propagation of RIP messages from the switch's private side.
- Configure the private side of the switch to RIP passive mode or active mode. In passive mode the switch listens for RIP messages and updates its routing table, but it does not advertise RIP routes. In active mode the switch updates its routing tables and also advertises its routes to other routers.

The interface filters setting can affect the behavior of routing protocols. For example, RIP uses UDP as its transport mechanism and if the interface filters are set to deny UDP, then RIP advertisements are not sent or received.

Protecting against routing loops

A routing loop can occur when two or more routers continuously forward the same packet to each other until the hop count goes to infinity, the packets time-to-live counter expires, or the network goes down. Loops typically occur when a new router is added to the network or when a router in an existing network goes away and the remaining routers must recalculate routes. A loop detection protocol can prevent a routing loop and can speed up convergence while the situation corrects itself. The switch supports the following methods used by RIP for minimizing loops and for speeding up the convergence that is caused by the normal correction of a loop:

- Split horizon, where the switch does not send routes that it learns from a neighboring router back to that same neighbor.

- Split horizon with poison reverse, where the switch *does* send back the routes that it learns from a neighboring router, but it sets the metrics for the connection to infinity.
- Triggered Updates, where an update is sent almost immediately after a routing change has been made on the switch. This is in contrast to the default RIP method, in which routes are updated at regular intervals.

About OSPF

OSPF is a link-state routing protocol that maintains a database from which a routing table is constructed from the shortest path, using a minimum of routing protocol traffic. It provides a high functionality open protocol that allows multiple vendor networks to communicate using the TCP/IP protocol family. Some of the benefits of OSPF are:

- Fast convergence
- Variable-length subnet masks (VLSM)
- Hierarchical segmentation
- Area routing to provide additional routing protection and a reduction in routing protocol traffic
- Authentication

Using OSPF

The switch sends OSPF broadcast/multicast messages at regular intervals. These messages contain information about routes that the switch can reach. Other routers on the network pick up these messages, and in turn update their routing tables and then send out route messages to their peer routers. The switch's OSPF support allows you to:

- Enable/disable propagation of OSPF messages from the switch's private side.
- OSPF passive mode or active mode. In passive mode the switch listens for OSPF messages and updates its routing table, but it does not advertise OSPF routes. In active mode the switch updates its routing tables and also advertises its routes to other routers.

The interface filters setting can affect the behavior of routing protocols. For example, OSPF uses UDP as its transport mechanism and if the interface filters are set to deny UDP, then OSPF advertisements are not sent or received.

Advanced routing and firewall keys

For OSPF to function properly, you must install the Advanced Routing Key. The Firewall Key is optional and required only when you want to enable the redistribution capabilities of RIP and OSPF. If you configure any OSPF policy without installing the applicable key, the configured policy will not be enforced.

About VRRP

Virtual Router Redundancy Protocol (VRRP) is one method you can use to configure the switch to maintain a state of High Availability. VRRP is a standard protocol that handles private interface failures. VRRP targets hosts that are configured with static next-hop routing addresses or default gateways. It provides a means of rerouting traffic in the event of a system/interface failure.

VRRP is managed as two separate parts. The first part handles those configuration parameters that must be the same between all switches that make up a Virtual Route (VR).

Use of an external LDAP server makes it easier to configure VRRP because it provides a common location in which information about each switch in the system can be maintained. Use of an external LDAP server enables each switch to see the settings of other switches on the system. Configuration of VRRP requires that Virtual Router IDs (VRIDs) are agreed to by all participating switches.

An external LDAP server is not a requirement. If the internal LDAP server is being used, the various switches must have these parameters configured the same and the responsibility for doing so lies with the administrator.

The second part of VRRP configuration is the information that is specific to a switch. This is information that is related to an interface and the role that the interface plays in VRRP (master or backup). This information is kept in the normal configuration file that is stored on the switch.

Configuring VRRP

The following steps describe how to configure your switch for VRRP.

- 1** Go to the Routing→VRRP screen, enter the IP address, and click Create.
- 2** On the VRRP→Edit VRRP IP Address screen, enter a decimal value between 1-255 for the VRID. This number must be unique to the LAN segment running VRRP.
- 3** In the Advertise Interval edit box, enter the interval at which the VR advertises its virtual MAC address. The range is 1-255 seconds and the default is 1 second.
- 4** Select None or Simple as the Authentication Type for this VR. None means that VRRP protocol exchanges are not authenticated and Simple means they are authenticated by a simple text password.
- 5** Enter up to 8 characters of text for the authentication string and confirm it.
- 6** You can optionally set up Master Delay Mode to control when a switch takes mastership of an address it owns.
- 7** Go to the Routing→Interfaces screen and check Enabled as the Configured State.
- 8** In the Backed up Address box, add all of the interfaces that you want to back up and click Add.
- 9** In the Master Status section, enable all interfaces that you want to be Master.
- 10** Click close.

Traffic patterns

The switch's enhanced routing feature provides a range of routing support. Table 16 shows the traffic patterns that are supported and lists which part of the routing table the traffic uses. This table also indicates whether the traffic goes through the switch's firewall.

Table 16 Traffic patterns

Traffic	Which routes are used when no firewall is enabled?	Which routes are used when a firewall is used?	Does it go through firewall when a firewall is used?
VPN routing and enhanced routing			
From tunnel: • tunnel→private • tunnel→tunnel • tunnel→public *	Private only	• Uses all routes. • Private default has precedence over public default.	• Yes if it is to a public interface • No for all other cases
From private interface: • private→private * • private→tunnel • private→public *	Private only	• Uses all routes. • Public default has precedence over private default.	• No if it is to a tunnel interface • Yes for all other cases
From public interface: • public→private * • public→tunnel * • public→public *	Traffic is dropped.	• Uses all routes. • Private default has precedence over public default.	• Yes
Services routing			
Inbound destined for switch: • private→Switch • tunnel→Switch • public→Switch	Does not use routing table.	Does not use routing table.	• No if it is from a tunnel • No if it is a tunnel protocol • No if it is to the management address • Yes for all other cases
Outbound from switch private interface address	Private only	• Uses all routes. • Public default has precedence over private default.	• No if it is a tunnel protocol • No if it is from management address • Yes for all other cases
Outbound from switch public interface address	Public only	• Uses all routes. • Private default has precedence over public default.	• No if it is a tunnel protocol • Yes for all other cases

* Indicates traffic goes through the firewall.

Client address redistribution

When a client initiates a user tunnel, the switch assigns an inner address to the client. Sources for these addresses can be:

- A predefined address pool in the switch with an address range that belongs to a locally attached private network
- A predefined address pool in the switch with an address range that does not belong to any locally attached private network
- A static address configured in the switch
- A RADIUS or DHCP address
- A client-supplied private address

If addresses from any of these sources do not belong to any of the locally attached switch networks, you must enable client address redistribution to ensure that these addresses are advertised in the dynamic route updates sent out by the switch. Client address redistribution uses a route type called a Utunnel. Utunnel routes can be either host or network routes.

When client address redistribution is active, the switch creates and advertises a user tunnel host route whenever a client tunnel is created using an inner address that does not belong to a locally attached network. When the tunnel is taken down, the corresponding host route is deleted.


When inner addresses are being allocated from an address pool with a range that does not belong to a locally attached network, you can reduce the number of entries in the route table and the route redistribution overhead by using the summarization option. Summarization creates and advertises a single Utunnel network route covering the address pool range when a client tunnel is created using an inner address from this address pool. The network route remains in the route table until the last tunnel using an inner address from this address pool is taken down.



Note: The maximum number of Utunnel routes cannot exceed the maximum number of client tunnels supported by the corresponding hardware platform. The default value is 200.

Figure 40 shows an example of client address redistribution where the client has an inner address that is not within the local subnet of the private network. The switch creates a Utunnel route that is then propagated over the network. This allows the router on the private network to recognize the 200.168.1.100 address and route responses back to it properly.

Figure 40 Client address redistribution



If you enable summarization, the switch identifies the subnet from the address pools where this address belongs and inserts a user tunnel network route for this subnet into the Route Table Manager. Summarization is disabled by default.


Enabling summarization is useful for large networks where optimization reduces the number of Utunnel host entries in the RTM. However, when using summarization, you could potentially have connection problems if the subnets of the address ranges span multiple switches. If you have two switches assigning addresses that belong to the same IP subnet, you should not use the summarization option.

For example, in Figure 41 Switch A has an address range of 200.168.1.100 to 200.168.1.120 and Switch B has an address range of 200.168.1.150 to 200.168.1.170. Both of these ranges are part of class C subnet 200.168.1.x/24. Client 1 logs in to Switch A and Client 2 logs in to Switch B. Both clients have

inner addresses that are not within the local subnet of the private network, but are in the same IP subnet. Switch A and Switch B running client address redistribution create Utunnel host routes. These routes are propagated over the network. The router on the private network recognizes addresses 200.168.1.100 and 200.168.1.150 and can route responses back to them through the designated switch.

If you enable summarization on both switches, both switches will advertise routes to 200.168.1.x/24. Router R will use one of these routes, causing either Client 1 or Client 2 to have communication problems.

Figure 41 Summarization for Client Address Redistribution



The Routing Table Manager handles Utunnel routes similarly to other route types (RIP or OSPF). You can view Utunnel routes using the Routing→Route Table Manager screen. The route policy service handles redistribution (advertisement) of utunnel routes similarly to redistribution of other route types.

To configure client address redistribution, go to the Routing→Client-Addr-Dis screen.

Configuring firewalls

The switch includes integrated firewall solutions that are designed to meet the needs of a variety of customers. The switch provides a choice of firewall solutions:

- Contivity Stateful Firewall
- Contivity Tunnel Filter
- Check Point FireWall-1.

Contivity Stateful Firewall

With the addition of the Contivity Stateful Firewall, the switch can perform a variety of secure routing functions, depending upon how you set up the switch's routing capabilities. For example, you can configure the switch to securely route non-tunneled traffic from its private interface, through the firewall, and out its public interface. This configuration would enable users on the switch's private network to access the Internet without requiring a separate, dedicated router.

By using stateful inspection, the Contivity Stateful Firewall provides a high level of security, the fastest runtime, and the flexibility to define the rules to fit your environment. The firewall delivers full firewall capabilities assuring the highest level of network security. To do this, the firewall examines packets in both incoming and outgoing directions running against a common security policy. All service rules are interpreted on IP conversations (not packets) and are fully stateful. Security rules do not filter packets directly, but the firewall services determine how to process them based on the security policy defined.

The firewall provides a user interface to help you determine the appropriate rules for your network. The Contivity Stateful Firewall achieves optimum performance as a result of advanced memory management techniques and optimized packet inspection. For further information on the Contivity Stateful Firewall, see *Managing the Contivity Stateful Firewall*.

Contivity Tunnel Filter

For many customers, the Contivity Tunnel Filter provides a cost-effective level of protection. Contivity Tunnel Filter can only be disabled when Contivity Stateful Firewall is enabled. You typically enable it while you are migrating to the Contivity Stateful Firewall.

Check Point FireWall-1

For customers who have Check Point FireWall-1 from Check Point Technologies, LTD, the switch offers an optional, separately licensed integrated firewall. It is transparent to both users and applications that access the switch. The Check Point FireWall-1 feature is an integrated, fully functional FireWall-1 (Version 4.0). To use the integrated FireWall-1, you must obtain a license for the firewall Inspection Module from Check Point Software Technologies, Ltd. You must also have a licensed version of the Check Point FireWall-1 Management Server. For further information on Check Point FireWall-1, see *Installing Check Point FireWall-1 on the Contivity Extranet Switch*.

Enabling your switch's firewall support

To enable your switch's firewall support, select one of the following choices on the Services→Firewall screen and click on the OK button. You are then prompted to reboot your switch. You can select only one firewall choice at any time.

- Contivity Stateful Firewall enables the switch's Contivity Stateful Firewall feature. When you enable the Contivity Stateful Firewall, you can run any combination of the following:
 - Contivity Stateful Firewall
 - Contivity Interface Filter
 - Interface NAT
 - Anti-spoofing
- Check Point FireWall-1 enables the switch's integrated Check Point FireWall-1 software.
- No Firewall disables all firewall features on the switch. In this configuration, the switch performs VPN routing only.

Confirming and saving your firewall selection

After you change your Firewall selection on the Services→Firewall screen and click on the OK button, you are prompted to confirm your selection. If you selected No Firewall, click on the OK button.

If you selected either the Contivity Stateful Firewall or the Check Point FireWall-1, you must restart your switch before the firewall becomes active. For these selections, you are given the following choices:

- Click on the OK button to reboot your switch now and use your new settings. The Admin→Shutdown screen appears. Complete the standard System Shutdown procedure to restart your switch.
- A warning message is continually displayed reminding you that you have made firewall changes and that you must reboot your switch to use the new settings.
- Click on the Cancel button to return to the Firewall screen without saving any of your changes and continue using your current firewall settings.

After you enable the firewall support, you must configure the specified firewall, if you have not already done so:

- For Contivity Stateful Firewall support, see *Managing the Contivity Stateful Firewall*.
- For Check Point FireWall-1 support, *Installing Check Point FireWall-1 on the Contivity Extranet Switch*.

Configuring IPX

The Internetwork Packet Exchange (IPX) protocol is the Novell adaptation of the Xerox Networking System (XNS) protocol. IPX has the following characteristics:

- It is a connectionless datagram delivery protocol. A datagram is a unit of data that contains all of the addressing information required for it to be delivered to its destination.
- It does not guarantee the delivery of packets. Higher-level protocols assume the responsibility for reliability.

IPX is the network-layer routing protocol used in the Novell NetWare environment. The primary tasks of IPX are addressing, routing, and switching information packets from one location to another on a network. In a LAN-based client the network interface card (NIC) provides network node addressing; in a tunneled environment, the switch provides the network node addressing.

Network addresses form the basis of the IPX internetwork addressing scheme for sending packets between network segments. Every network segment of an internetwork is assigned a unique network address by which routers forward packets to their final destination network. On the switch, all public interfaces are treated as a single network segment with a unique network address. A network address in the NetWare environment consists of eight hexadecimal characters. In the following example, 0x indicates that this is a hexadecimal number, and *n* is any hexadecimal character.

0xxxxxxxxx

Socket numbers are the basis for an IPX *intranode address* (the address of an individual entity within a node). They allow a process (for example, IPX Routing Information Protocol [RIP] and Service Access Points [SAP]) to distinguish itself to IPX. To be able to communicate on the network, the process must request a socket number. Any packets IPX receives addressed to that socket are then passed on to the process within the node.

The switch uses IPX RIP and SAP to dynamically learn and advertise IPX routes and services. The switch assigns IPX addresses to tunneled clients; remote users cannot configure the IPX tunnel address for their systems.

The switch does not forward IPX packets from a private nontunneled LAN to another private nontunneled LAN, nor does it propagate routing or server tables from a private nontunneled LAN to another private nontunneled LAN.

The switch supports IPX by encapsulating IPX traffic within IP tunnels over PPTP, L2TP, and L2F. Note that the switch's IPX support is not available for the IPSec tunneling protocol.

IPX client

On the PPTP client (for example, Microsoft Dial-Up Networking*), you must enable the dial-up networking IPX option. Enabling the IPX option allows you to tunnel using IPX, IP, or IPX and IP according to the dial-up networking selections.

Windows 95 and Windows 98

If you are running Windows 95 or Windows 98, you must load the intraNetWare* client, which is available from the Novell Web site:

<http://www.novell.com>



Note: The NetWare* client for Windows 95 and Windows 98 does not function properly; therefore, you must use the Novell intraNetWare client when using IPX with PPTP.

Windows NT

The NetWare client is already on Windows NT systems. You can use that or the Novell intraNetWare client, which you can access from the Novell Web site at <http://www.novell.com>.

IPX group configuration

IPX is disabled on a per group basis by default. Therefore, you must enable IPX for group users to access IPX. You enable IPX for group users from the Profiles→Groups→Edit→Connectivity screen.


Sample IPX VPN switch topology

In the following figure, the public interface IPX network address that the switch provides is 0000A100. As previously stated, regardless of the number of IPX public interfaces that are configured on the switch, they all use the same IPX network address.

You must enable the private interfaces that you want to use for IPX traffic, and for each private interface you must configure the IPX network address and IPX frame type. The IPX network address that you configure must match the IPX network address for that LAN, and the IPX frame type must match the IPX frame type for that LAN.

For example, in [Figure 42](#), the private interface network address to the NetWare server is 00000B16 and the Frame Type is 802.3; similarly, the private interface network address to the Nortel Networks Router is 00000C22 and the Frame Type is SNAP.

Figure 42 IPX topology



Note: The Private LAN can also carry IP and IPX traffic simultaneously. The IP addresses are not shown in this figure.

Chapter 6

Configuring Servers

This chapter describes how to configure the following authentication servers for users who are tunneling into the switch:


- External LDAP is a standard protocol for Internet directory services that is based on directory entries. A directory service is a central repository of user information.
- Internal LDAP stores the Group and User Profiles on the internal server of the switch.
- External RADIUS is a distributed security system that uses an authentication server to verify dial-up connection attributes and authenticate connections.
- RADIUS accounting logs user sessions with RADIUS-style records containing detailed connection statistics.

The switch can function as a simple RADIUS server.

Authentication services

The switch supports LDAP and RADIUS Authentication Servers. The switch always attempts to authenticate a remote user against the LDAP database. If a user ID and password are found, the switch uses the attributes that are defined for that user's group. The switch can also authenticate against a RADIUS database. When using RADIUS for authentication you can create groups to take advantage of different profiles, or you can simply assign all RADIUS users into a single default group.

[Figure 43](#) shows a switch and authentication servers.

Figure 43 Contivity VPN Switch and authentication servers

The user ID (UID) is checked against the LDAP database. If the UID is found in the LDAP database, the user is assigned to a group and acquires that group's attributes. Next, the password is checked, and if it is correct, the switch allows a tunnel to be formed.


If the UID is not in the LDAP database, then the UID and password are checked against the RADIUS database. If the UID and password are correct, the switch checks to see if the RADIUS server returned a Class Attribute. The RADIUS Class Attribute is treated as an LDAP group name. If a RADIUS Class Attribute is returned, and it names an existing LDAP group, the switch applies the attributes of this group to this user's session, and forms a tunnel. If the group name does not exist, the user is given the RADIUS default group's attributes. If the UID and password are incorrect, the switch rejects the user request.

IPSec behaves the same as a PPTP session; the RADIUS server defines the group for that user after authentication using the Class Attribute Group Identifier. The only difference between IPSec and PPTP is that in the event the RADIUS server does not return a class attribute, the group associated with the IPSec Group ID is

used instead of the RADIUS default group. You configure the IPSec Group ID in the Authentication section of the Profiles→Groups→Edit→Configure IPSec screen. You configure the PPTP default group on the Servers→RADIUS Auth screen, RADIUS Users Obtain Default Settings from the Group option.

The group that the user is bound to must allow the authentication method that is used when the session is started.

Figure 44 illustrates the steps in Nortel Networks user validation.

Figure 44 Authentication server validation flowchart


LDAP, RADIUS, and the IPsec client

Authentication using the Nortel Networks IPsec client provides several options for remote users connecting through a switch. These include:

- UID and Password authentication against an LDAP database
- Group Password Authentication using a RADIUS server
- Token Security methods (Security Dynamics SecurID* and AXENT Omniguard/Defender*)

Figure 45 shows IPsec client options.

Figure 45 IPsec client authentication options



Each authentication option has the following positive attributes:

- Diffie-Hellman key exchange (ISAKMP/Oakley Aggressive Mode) to build the Security Association (SA).
- The user name and the password are never transmitted in the clear; a cryptographic hash function (SHA-1) is used to protect the user's identity.

- Mutual authentication between the client and the switch using a keyed hash algorithm (HMAC).
- Protection against authentication replay attacks through the use of session “cookies.”

LDAP database authentication

The IPsec client and the switch support authentication against an LDAP database. This database can reside internally in the switch or it can reside on a server. You can configure the switch to authenticate the remote access user with a simple ID and password. In this mode, the user ID is protected from network snooping by hashing the user name with a standard one way function (SHA-1). Mutual authentication is achieved by each knowing the password hash and the use of data “cookies” by each side to prevent replay attacks.

To enable IPsec authentication against the LDAP database:

- 1 Select this option in at least one group profile from the Profiles→Groups→IPsec: Edit screen.
- 2 Select User name and Password as one of the Authentication options.
- 3 Add users to the group for authentication from the Profiles→Users screen; select Add User for the appropriate group.

RADIUS-based authentication

The switch supports authentication against a RADIUS server. This server can reside on either a private or public network that is connected to the switch. To enable RADIUS authentication, you must configure the switch with the RADIUS server host name, port number (typically 1645), and a shared secret. Access the switch management screen from the Servers→RADIUS Authentication screen.

The RADIUS Authentication screen also allows you to configure the type of authentication methods that are allowed to access the RADIUS server. There are five options of which only four are IPsec-related; AXENT, SecurID, CHAP, and PAP. MS-CHAP is available for PPTP tunnel users only (it is not applicable to IPsec tunneling applications).

If you are using token cards for authentication, you must select the appropriate technologies (AXENT and/or SecurID). If you want to use simple user IDs and passwords to authenticate against the RADIUS server you must select CHAP or PAP. Under no circumstances is the UID or password passed in the clear from the remote client or from the switch communicating with the RADIUS server. From a security point of view, there is no significant difference between using CHAP or PAP since the connection between the switch and the RADIUS server is protected by encryption. PAP Authentication consumes fewer instructions during the authentication process, which is a minor consideration.

When using RADIUS-based authentication, the IPSec client and the switch require a second set of credentials that are used for mutual authentication. These credentials are referred to as the Group ID and Group Password.

The remote access client information is documented in the Contivity VPN Client online Help. On the IPSec client side, the remote user must select Options→Authentication Options, and then click on: User Group Security Authentication. Next, the remote user enters the Group ID and Group Password that you provide. And lastly, they must select one of these options:

- Challenge Response Token
- Response Only Token
- Group Password Authentication

To complete the RADIUS setup, you must configure at least one group profile for RADIUS users. In this profile you need to enter the Group ID, password, and the allowed Group Authentication Options. You can configure the group profile from the Profiles→Groups→IPSec: Edit screen.

- 1** Set up and test the operation of the RADIUS server with ACE and/or Defender servers, depending on the type of Token Security you want. You should do this before attempting authentication by an IPSec client to verify that everything on this side of the network is operating properly.
- 2** Identify and create the groups for authenticating token users, and supply the Group ID and Password to all users doing either Token Card or Group Password Authentication. Note that AXENT and SecurID users are created and maintained in their respective servers, not in the switch. Add the groups in the Profiles→Groups→IPSec screen.

- 3 Define the RADIUS server configuration settings for Token Security.
- 4 Define the Tunnels settings for IPsec. Add a RADIUS server, if necessary.

RADIUS authentication servers

RADIUS is a distributed security system that verifies dial-up connection attributes and authenticates connections.

The RADIUS Authentication Servers screen allows you to configure up to three servers for remote authentication. It is imperative that the RADIUS servers contain the same user data. The alternative RADIUS servers are used only when no response is received from the Primary RADIUS server.

Most RADIUS servers support CHAP and PAP authentication, and some can now support MS-CHAP (Funk, for example). If you require PPTP-encrypted tunnels and RADIUS authentication, then you must use a RADIUS server that supports MS-CHAP. The alternative is to use an LDAP server for PPTP authentication.

RADIUS authentication class attribute values

The following illustration shows the relationship between RADIUS authentication class attribute values for switch users. In [Figure 46](#), C is the class attribute for country, and OU is the class attribute for organizational unit.

Figure 46 RADIUS authentication class attribute values




Table 17 shows sample details that you enter into your RADIUS server.

Table 17 RADIUS example details

User ID	Class attribute value	Assigned group
Lee Madison	ou=New Products, ou=Research and Development	New Products
Julie Lane	None	Default
Bill Sullivan	ou=Staff	Default (since ou=Staff does not exist)

The RADIUS server uses the Class Attribute Value to associate the user ID with a group in the LDAP database.

RADIUS-Assigned Framed-IP-Address attribute

You must configure a RADIUS-Assigned Framed-IP-Address attribute (8) on the RADIUS server for the ID being authenticated by the switch. If the option Allow Static Addresses (Profiles→Groups→Edit→Connectivity screen) is enabled for the assigned group, then the returned IP address is used for the tunnel session. Otherwise, an IP pool address is assigned.



Note: Only a single IP address is returned by the RADIUS server; therefore, only one active tunnel connection is permitted per user ID.

Configuring the switch for RADIUS

The following procedures describe how to configure the switch to interoperate with a RADIUS server while using either IPsec and PPTP.

IPsec and RADIUS

To configure IPsec and RADIUS on the switch:

- 1 Go to the Servers→Radius Auth(entication) screen and select Enable Access to RADIUS Authentication.
- 2 Enable an authentication method from the Server Supported Authentication Options.
- 3 Enable a server and enter the server's Host Name or IP address, the Interface type, Port Number (1645), and Secret. Click OK.
- 4 Go to the Services→IPsec screen and use the Add RADIUS button to add a RADIUS server to the Authentication Order table.
- 5 Go to the Profiles→Groups screen and either add or select the group that you want to be the default group for RADIUS users (this is the group a user is assigned to if the RADIUS server does not send back a class attribute).

- 6 Next, go to the Profiles→Groups→IPSec Configure screen. In the Authentication area, select the Configure button.
- 7 On the next screen, the Authentication method for the Group is already selected. Enter the Group ID and Group Password. Consider using the LDAP group name as the default group, because you must remember a default name once you enter it. If your RADIUS server returns a class attribute, ensure that the Authentication method is enabled for that group. However, you do not need a Group ID and Group Password for the group that is being returned as a class attribute.

To configure IPSec and RADIUS on the client:

- 1 In the Contivity VPN Client GUI, go to Options→Authentication Options, select Group Security Authentication, and enter the Group ID and Group Password.
- 2 Select one of the Group Authentication Options.
- 3 Click OK.

PPTP and RADIUS

To configure PPTP and RADIUS on the switch:

- 1 Go to the Servers→Radius Auth(entication) screen and select Enable Access to RADIUS Authentication.
- 2 Enable an Authentication method.

If a valid class attribute is not returned, then PPTP users are placed in the default group as configured on the Servers→RADIUS Auth(entication) screen.



Note: Everything about the authentication type must match; for example, if you send an encrypted password then MS-CHAP must be enabled on the RADIUS authentication screen and the RADIUS server must support MS-CHAP.

AXENT software tokens and the IPSec client

The following steps enable the switch to work with AXENT software tokens and AXENT's RADIUS server using the IPSec client. You want the switch to look and appear like a network access server (NAS) to the AXENT RADIUS server. You must have AXENT RADIUS Server version 3.0 or later.

- 1 Go to RADIUS→Authorization Profiles on the AXENT Defender user interface.
- 2 Create a profile for Contivity VPN Switch-authenticated users that has a Service Type of Authenticate Only and a framed protocol of PPP.
- 3 Associate this profile with all switch users.
- 4 Make sure you have the proper class configuration for dial-in security for each user profile. That is, make sure the associated class allows tokens (for example, DHT Required), and that RADIUS is an allowed authorization type.
- 5 Enable the switch as a valid Agent and set the agent key properly on the DMS Remote Manager screen.
- 6 Perform a DB Sync or RADIUS Sync from the DSS Remote Manager screen after making these changes.

The AXENT RADIUS server can operate with different levels of logging as configured from the command line. This logging is a useful debugging tool. Generally, problems occur when the RADIUS server cannot communicate correctly with the DSS server. This is apparent in the log file.

LDAP servers

The Lightweight Directory Access Protocol (LDAP) is a standard protocol for Internet directory services that is based on directory entries. A directory service is a central repository of user information; for example:

- Groups
- Users
- Filters
- Services

An entry is a collection of attributes that has a distinguished name (DN), which refers to the entry unambiguously. Each entry attribute has a type and one or more values. Types are typically mnemonic strings; for example, cn represents common name and mail represents e-mail address. The values depend on the attribute type. For example, a mail attribute value might resemble `jchirac@elyseepalace.france.gov`.

LDAP directory entries are arranged in a hierarchical tree-like structure that reflects political, geographic, and organizational boundaries. Country entries appear at the tree top. The next entries represent states or national organizations. The third-branch entries represent people, organizations, servers, files, or any other readable database entry. LDAP allows you to read, search, add, and remove information from the centralized database.

The switch uses an LDAP server to centrally store remote access profiles and corporate networking details such as the addressing mechanism; for example, group attributes including hours of access, filters, and authentication servers. The switch queries the LDAP server for access information when a user establishes a tunnel connection. The LDAP query can be serviced locally by the internal LDAP server; or it can be redirected to an external LDAP server, such as the Netscape Directory Server,* Microsoft Active Directory* Service, or Novell Netware* Directory Services.

The switch's internal LDAP server is designed to operate only as an internal server and does not respond to external queries. Therefore, two or more switches cannot share the same internal LDAP database. To allow sharing between switches, and to take full advantage of LDAP-based directory service replication and centralization, you should use a dedicated directory services offering from one of the above vendors.

The switch synchronizes its cache every 15 minutes. For example, if you delete a user from an external LDAP database it can take up to 15 minutes before all of the switches recognize the change. Additionally, the LDAP server's status is recorded in the Event Log every 15 minutes.

To configure internal LDAP, go to `servers→LDAP`.

Using a Netscape directory server

Nortel Networks is currently providing Netscape Directory Services Version 4.0 on the Contivity VPN Switch CD.

You cannot currently perform LDAP authentication over a Branch Office connection. To bring up a Branch Office connection, you must have a local LDAP server (either an internal server or an external server on the private network) to authenticate the remote site.

If you are using an external LDAP server, it must be on the private network. To install the Netscape Directory Server Software:

- 1** Place the Contivity VPN Switch CD in the Windows LDAP server CD-ROM drive. Locate the LDAP/Netscape folder on the CD.
- 2** Unzip the Netscape file, and double-click on the Setup icon to start the Netscape Directory Server installation.
- 3** Follow the installation instructions, selecting the defaults. The Netscape server software takes approximately 50 MB of disk space after the installation is complete, plus space for the user database itself. Supply an administrator user name and password when you are prompted to do so. This is the user name and password that is required to run the Netscape Console administration tool on the LDAP server.

At the end of the installation, the installation procedure prompts you to launch the Netscape Console administration tool. This is a Java-based management tool.

Quick start

The following instructions describe setting up a Netscape Directory Server (NSDS, Versions 3.x and 4.0) to function as an authentication server for your switch. Generally, these steps pertain to all users. Some of the steps pertain only to users who have installed the Netscape Directory Server from a source other than the Nortel Networks CD.

The default configuration directory location of a Netscape Directory Server 3.x is `/netscape/suitespot/slapd-<server_name>/config`. The default configuration directory location of a Netscape Directory Server 4.0 is `/netscape/server4/slapd-<server_name>/config`. Unless otherwise stated, all files to be edited are in these default configuration directories. The main configuration file in this directory is `slapd.conf`.

Follow these steps to connect your switch to a newly installed Netscape Directory Server. These instructions assume that Secure Socket Layer (SSL) is not being used.

For Version 3.1, from a source other than Nortel Networks CD installation only:

- 1 Stop the Netscape Directory Server. Edit `slapd.dynamic_ldbm.conf` as the configuration directory in the following instructions.
- 2 Using FTP, transfer the file `/system/slapd/config/bnes.idx` from the switch into the default configuration directory. Add the following line to the end of the Netscape file:

```
include "<config_directory>/bnes.idx"
```

where `<config_directory>` is the full path to the configuration directory on the Netscape Directory Server disk.

For all others:

- 1 Start the Netscape Directory Server.
- 2 Go to the Servers→LDAP screen and choose: switch to External Server.
- 3 Enter the value of the first Netscape 3.x `slapd.conf` file suffix or the Netscape 4.0 `slapd.ldbm.conf` file suffix attribute in the Base DN field.
- 4 Enter the value of the Netscape `slapd.conf` file `rootdn` attribute in the Master server Bind DN field. This is the same DN that you entered as the *Unrestricted User* when you created your Netscape Directory Server.
- 5 Enter the same Bind Password and Confirm Password as you entered for the *Unrestricted User* when you created the Netscape Directory Server.
- 6 Click OK and go to the Profiles→Users screen. After a brief delay you see a screen with the group `/Base`.

You are now connected to the external LDAP server and your skeletal database has been created. You can add users, groups, filters, and hours of access records through the Profiles screens. To connect other switches to the same Netscape LDAP server, repeat steps 2 through 6.

Advanced configuration

This section contains details for modifying the configuration of your Netscape Directory Server. This section details the procedure for modifying the server's schema attribute index list and SSL configuration, and then configuring the switch to connect to it. Netscape Directory Server 4.0 for Windows NT is provided on the Nortel Networks CD.

You must first disable the UID uniqueness plug-in through the Netscape console. This is located in the Netscape Directory Server Console Configuration tab. And you must restart the Netscape Directory Server after disabling the plug-in.

For an NSDS 4.0 server, no other configuration changes are required as long as the switch binds to the server as an appropriately privileged user. The switch attempts to modify the schema and attribute indices of the server after connecting to it. To do this, the switch must have write access to the schema entry (cn=schema) and the database configuration entry (cn=index, cn=config, cn=ldbm). If you are using the root DN as your bind DN, the switch already has the required access privileges to write to the schema and database configuration entries.

Schema/attribute index cleanup

If you have NSDS 3.x or earlier installed, its switch schema modifications must be updated manually. All attributes, objectclasses and attribute indexes prefixed with newoak must be removed. These attributes and objectclasses are contained in the files config/slapd.at.conf and config/slapd.oc.conf. The indexed attributes are contained in config/slapd.dynamic_ldbm.conf. You must remove all lines that contain an attribute or objectclass that is prefixed by newoak from these files.

Schema/attribute index modifications

You must manually modify the NSDS 4.0 schema only if the switch bind DN (cited earlier) that is used to access the server does not have privileges to modify the schema or database index entries.

To manually modify the schema or database index entries:

- 1 FTP the following files from the /system/slapd/ldif directory on the switch to any directory on your NSDS file system:
 - bnes.idx
 - bnes.at
 - bnes.oc
- 2 Add the following line to the end of the config/slapd.dynamic_ldbm.conf (NSDS 3.x) or config/slapd.ldbm.conf file (NSDS 4.0), where *<config_directory>* is the full directory path used in step 1:

```
include "<config_directory>/bnes.idx"
```
- 3 Add the following line to the end of the config/slapd.user_at.conf file, where *<config_directory>* is the full directory path used in step 1:

```
include "<config_directory>/bnes.at"
```
- 4 Add the following line to the end of the config/slapd.user_oc.conf file, where *<config_directory>* is the full directory path used in step 1:

```
include "<config_directory>/bnes.oc"
```
- 5 Restart the NSDS server.

SSL

You must configure the Netscape Directory Server with compatible encryption types and you must have a certificate signed by a Certificate Authority that is trusted by the switch to use Secure Socket Layer (SSL). You enable SSL and configure encryption types on the Netscape Server Administration management screens. You install certificates and keys through the Netscape General Administration screens.

To configure the switch:

- 1 Start the Netscape Directory Server.

- 2 Go to the Servers→LDAP screen and choose switch to External Server.
- 3 Enter the Netscape Directory Server DNS host name or IP address in the Host Name field.
- 4 Change the port or SSL port number, as necessary, to the port selected when the Netscape Directory Server was configured. Note that the SSL port must be selected in order to have the switch use SSL to connect to the Netscape Directory Server. Also, the Netscape Directory Server's certificate must be signed by a trusted Certificate Authority and must support at least one compatible encryption type in order to make a successful connection.
- 5 Enter the distinguished name, in the Base DN field, of the entry in your Netscape Directory Server that is serving as the base entry for the switch.

In previous switch releases, this entry could not exist. Now, this entry can exist and the switch modifies its contents. The parent entry must always exist. For example, if the Base DN is:

```
ou=Remote Access Users, ou=East, o=MyCompany, c=US
```

The organizational unit, Remote Access Users, can exist but does not have to exist, but its parent organizational unit, East, must exist.

For a new Netscape Directory Server 3.x, you must specify the following, since no entries yet exist in the database, as the Base DN (assuming this is also the value of the suffix attribute):

```
o=MyCompany, c=US
```

An entry similar to the preceding one appears in the 4.x database automatically. This causes the database and Base DN of o=MyCompany, c=US to be created on the Netscape Directory Server. The Base DN is used as the Base group for the switch.

- 6 In the Bind DN field for the Master server, enter the DN of a user that has read and write access to the Base DN and its subentries. For the Slave servers, this DN only needs read access to the Base DN and its subentries.
- 7 In the Bind Password and Confirm Password fields, enter the password that corresponds to the Base DN.
- 8 Click on OK. After a short delay you see either “operational” or “error” in the server status column. If the status is error, something prevented the switch from successfully converting to the server. Possibilities are:

- The IP address, port bind DN, or password is incorrect.
 - The schema on the server either could not be updated (NSDS 4.0) or is incorrect.
 - The UID uniqueness plug-in is enabled (NSDS 4.0).
 - The base DN is not an organizational unit (ou) or an organization (o).
- 9 If the status is operational, you are now connected to the external LDAP server and a skeletal database (contained in the switch's file `/system/slaped/ldif/template.ldf`) has been created. You can now add users, groups, filters, and hours of access records through the Profiles screens.

Exporting internal LDAP to Netscape directory server

This procedure describes exporting a database from the switch's Internal LDAP server to a Netscape Directory Server. The default LDIF directory location of a Netscape Directory Server 3.x:

```
/netscape/suitespot/slaped-<server_name>/ldif
```

The default LDIF directory location of a Netscape Directory Server 4.0:

```
/netscape/server4/slaped-<server_name>/ldif
```

The Netscape Directory Server must first be prepared for additional information. To export a database from the switch's Internal LDAP server to a Netscape Directory Server:

- 1 Click Servers→LDAP.
- 2 Enter a name in the Backup to File field.
- 3 Click Backup Now. Wait for the backup to complete (this can take several minutes if the database is large). You can monitor the backup status on the Status→Health Check screen. The status of the Internal LDAP Server is “backup in progress” during the backup and reverts to “Internal Server is down” when the backup is complete.
- 4 Using FTP, transfer the following file from the switch to the Netscape LDIF directory:

```
switch/system/slaped/ldif/<backup_file_name>
```

- 5 Edit the file to change all occurrences of the string `o=Nortel Networks, c=US` to the value of your Directory Suffix (for example, `o=MyCompany, c=US`). This might be a subentry of your Directory Suffix as long as the parent of this entry exists. For instance, if your database contains the entry:

```
ou=East, ou=Sales, o=MyCompany, c=US
```

You might want to replace `o=Nortel Networks, c=US` with:

```
ou=Remote Access Users, ou=East, ou=Sales, o=MyCompany, c=US
```

In this case, you also need to change two other lines in the base entry of the exported LDIF file. The base entry is the entry containing the line:

```
o:Nortel Networks
```

Change this line to:

```
ou:Remote Access Users
```

In the same entry, change the line:

```
objectclass: organization
```

to:

```
objectclass: organizationalUnit
```

- 6 Restart the Netscape Directory Server (if necessary).
- 7 To add all entries in the LDIF file to the Netscape Directory Server database, go to the LDIF directory and enter `../bin/slaped/server/ldapmodify -D<bind_DN> -w<password> -a -f <ldif_file_name>` or use the graphical interface provided by Netscape to load the LDIF file.
- 8 Change the switch's LDAP server from Internal to External and configure it to use the Netscape Directory Server. Certificate validity check consists of the following:
 - Are the certificate Start and End Dates valid?
 - Is the basic constraint extension present; if so, does it indicate an end entity?
 - Does the certificate's key usage indicate digital signature?
 - Is the certificate signature by the CA valid?

- 9 Use the certificate to verify the signature on the ISAKMP packet.

RADIUS accounting configuration

The RADIUS Accounting configuration screen allows you to specify how your switch saves RADIUS Accounting results. By default, the results are stored locally. You can optionally also save the RADIUS Accounting information to a remote RADIUS Server.



Note: If you set the date ahead and then set it back, external RADIUS accounting no longer works.

To configure RADIUS accounting, go to Servers→Radius Acct.

Remote user IP address pool

Remote access users who are using tunneling protocols require two IP addresses to form packets. The addresses are normally referred to as outer and inner addresses. The outer address, or public address, is visible when packets are traveling through the public data networks (PDNs). This address is negotiated between the client and the ISP to which it is connected. The switch does not have control of this address.

The inner IP address is the one that eventually appears on the private network when the outer layers of the packet are removed. Therefore, this address must lie within the private network address space. The switch provides the remote user with the inner IP address during tunnel setup. This address can come from an internal address pool, from an external DHCP server(s), or from a RADIUS server.

The switch assigns the inner IP address from one of several sources, using the following order:

- 1 User-specified (excluding IPSec)
- 2 Static address, either the switch's LDAP database or RADIUS server

- 3 Local address pool, either the switch's internal address pool or the DHCP-acquired address pool

To configure the user IP address pool, go to Server→User IP address pool.

Configuring the switch as a RADIUS server

The RADIUS Service feature allows the Contivity VPN Switch to function as a simple RADIUS server. When the feature is enabled, the switch accepts RADIUS authentication requests from an external RADIUS client, then searches its LDAP database for the requested user. If the user is listed in the database and has the proper credentials, a positive response is returned to the requesting system. Otherwise the authentication request is rejected. Here are some key points about RADIUS Service that you should consider:

- RADIUS Service accepts standard RADIUS Access-Request packets. It then returns either a standard RADIUS Access-Accept or Access-Reject packet, based on the authentication result. The packet contains only a Reply-Message attribute. No other attributes are returned.
- RADIUS Service supports PAP, CHAP, and MS-CHAP authentication methods.

For users with multiple user accounts, RADIUS Service attempts to authenticate against each account type. If the given username/password matches any of the user's accounts, the authentication succeeds. The authentication is done in this order: PPTP, IPSec, L2F, L2TP.

Chapter 7

Using Certificates

The LDAP Server Secure Socket Layer (SSL) Encryption allows you to configure SSL encryption. SSL provides Internet security and privacy and ensures privacy between the switch and the external LDAP server. The SSL protocol negotiates encryption keys and authenticates the server before any data is exchanged. SSL maintains the transmission channels' security and integrity through encryption, authentication, and message authentication codes. SSL is considered very secure and supports the following encryption methods:

- RC4 128-bit MD5 allows clients to request RC4 128-bit MD5 encryption. This is the most secure method. The longer the encryption key, the more secure the encryption. US export law controls the export of 128-bit encryption keys.
- DES 56-bit SHA allows clients to request DES 56-bit SHA encryption. This is the mid-level encryption method, less secure than RC4-128 yet more secure than RC4-40.
- RC4 40-bit MD5 allows clients to request RC4 40-bit MD5 encryption. This is the least secure method of encryption.

SSL and digital certificates

The SSL Protocol can use digital certificates when establishing secure, authenticated connections between SSL clients and servers. Digital certificates are based on the X.509 certificate standard. Certificates are credentials created by a CA to assure a person's identity. Digital certificates verify that a specific public key belongs to a specific individual. CAs are either self-created (within an organization), or they are third-party companies or organizations.

The Contivity VPN Switch uses a digital certificate sent from an SSL-capable LDAP server to authenticate that server. For digital certificate authentication to succeed, a certificate from the authority certifying the LDAP server must be imported into the switch's certificate store. This type of certificate is often referred to as a CA Root certificate.

A single CA Root certificate can be used to certify the authenticity of multiple LDAP servers depending on the organization of your environment's certification hierarchy.

The SSL protocol enables you to establish a secure connection over an insecure channel. SSL is layered between TCP/IP and higher-level services such as HTTP, SNMP, LDAP, and Telnet. The switch only uses SSL with the LDAP protocol when accessing an SSL-capable LDAP server. The SSL security handshake negotiates security levels upon initialization of a TCP/IP connection. Additionally, SSL also negotiates an encryption algorithm and cryptographic keys before the application protocol transmits or receives any data.

The SSL protocol provides the following connection security properties:

- A private connection. Encryption happens after an initial handshake to define a secret key. Data is encrypted using DES or RC4.
- The peer's identity is authenticated using asymmetric, public key cryptography, such as RSA or the Digital Signal Standard (DSS).
- A reliable connection. Message transport includes an integrity check using a keyed MAC. MAC computations use secure hash functions (SHA and MD5).

These are the basic steps in an SSL digital certificate exchange:

- 1** The client (the switch, in this case) sends a connection request to the server.
- 2** The server sends the client a signed certificate containing the server's public key.
- 3** The client verifies that the certificate signer is on its acceptable Certificate Authority (CA) list.
- 4** The client generates a session key that is used for encryption and sends it to the server encrypted, using the server's public key.
- 5** The server uses its private key to decrypt the client-generated session key.


- 6 The client makes an LDAP request and the server provides an LDAP response that is encrypted using the exchanged session key. All subsequent communication with this LDAP server is also encrypted.

When you bring up an SSL connection, you receive a certificate in the network exchange that will be signed by a Certificate Authority. The switch contacts the LDAP server over a TCP connection to request a secure session. Along with other connection establishment information, the LDAP server sends a certificate to the switch.

The LDAP server can send a chain of certificates. Such a chain can contain CA certificates signed by other CAs. To verify the authenticity of the LDAP server's actual certificate (contained at the end of the chain), the switch examines the chain of CA certificates until it finds a CA certificate that has been signed by a certificate authority trusted by the switch (for example, OU=Secure Server Certificate Authority, Inc., O=RSA Data Security, C=US).

Using the trusted CA's certificate from the switch's certificate store, the switch verifies the chain of certificates starting with the CA certificate signed by the trusted CA verifying all certificates in the chain down to the final LDAP server's certificate. Verification consists of checking the validity date of the certificate as well as verifying the certificate's digital signature. [Figure 47](#) shows a sample certificate chain.

Figure 47 Sample certificate chain



For example, the LDAP server sends the certificate chain shown. The switch has CA Certificate B, for example, in its CA certificate store, as previously imported to the switch. By marking CA Certificate B as trusted, anything that is signed by B is trusted by the switch. For this example, since CA Certificate C has been signed by B, the switch trusts Certificate C if Certificate C's signature can be verified. Once Certificate C is trusted, the switch now trusts anything that Certificate C has signed. Since Certificate C signed the LDAP server's certificate, the switch now verifies the LDAP server's certificate and if it is verified, then the switch authenticates the LDAP server.

Once verified, the switch extracts the LDAP server's public key and encrypts a secret session key that it returns to the LDAP server. The LDAP server then decrypts this secret session key with its private key. Along with the original certificate chain, the LDAP server also sends a cryptographic vector, which informs the switch of the encryption options (RC4 or DES) the LDAP server is capable of performing.

Once the switch knows the encryption algorithm, and has its shared secret, both sides only exchange encrypted data over the LDAP connection. If there are no trusted authorities in the chain, then the switch rejects a connection request.


Tunnel certificates

You can use X.509 certificates to authenticate IPSec-based tunnel connections. The switch supports RSA digital signature authentication in the IPSec ISAKMP key management protocol. Remote users can authenticate themselves to the switch using a public key pair and a certificate as credentials. In addition, the switch uses its own key pair and certificate to authenticate the switch to the user.

The switch currently supports the Entrust product suite (Server: Entrust PKI 4.0 - Web Connector, VPN Connector, and Entrust Authority; Client: Entrust Entelligence 4.0, and VeriSign* OnSite*).

Using certificates for tunnel connections requires the creation of a public key infrastructure (PKI) to issue and manage certificates for remote users and switch servers. The switch software supports the Entrust PKI, Web Connector, VPN Connector, Entrust client software PKI components, and VeriSign OnSite service. The Entrust or VeriSign software issues certificates for users and the switches.

[Figure 48](#) shows the digital certificate interactions with the switch.

Figure 48 Digital certificate components

The PKI certificate infrastructure involves multiple components and multiple interactions when getting certificates, using certificates, and maintaining certificates.

Public Key Infrastructure (PKI)

An important decision you must make about the design of a PKI is how you provide CA services. You can use commercially available products from a vendor such as Entrust, where the CA would reside in your facility and be controlled and operated by you. Or, you can subscribe to a CA provider, such as VeriSign's OnSite service, where the CA would actually be operated by VeriSign from a remote location.

CA and X.509 certificates

The CA creates and revokes certificates. By publishing a certificate in its public repository, the CA certifies that it has issued the certificate to the named subscriber. The CA operates on information submitted to it by switch administrators and users of certificate client software, and it stores the results of its work in its LDAP-accessible Certificate Repository.

Certificate application, creation, and installation

Switch administrators who want to obtain certificates for their switches use the System→Certificates: Generate Certificate Request button to create a key pair, and to provide the public portion of the key pair to the CA in a PKCS #10-style certificate request. The certificate request is submitted to the CA by copying the request into the browser's clipboard (Edit→Copy) or by saving the request to a file. You then submit the request to the CA following the CA's instructions for handling a PKCS #10 request.

Entrust

Depending on which modules are installed as part of the Entrust PKI, the switch administrator must create a server certification request. This request must have a subject distinguished name with a common name that is equal to the Entrust Reference Number that is used to preauthorize the certificate issuance. The Entrust CA Administrator processes the request and creates the certificate. Depending on the Entrust product that is at the front end of the CA, you can retrieve the certificate directly using your Web browser or from a file storage location.

For the Entrust PKI with the Web Connector, you can use your Web browser to copy and paste the certificate request and response. With the Entrust VPN Connector, the certificate requests and response are transported using files. In all cases, you transfer the encoded response into your browser's cut-and-paste buffer, then import the certificate using the System→Certificates: Import Tunnel Certificate button.

VeriSign

In the default VeriSign OnSite configuration, the page on which certificate subscribers enroll is also hosted on a VeriSign Web server. The URL of this enrollment page is assigned during the registration of the OnSite account. For a private Client OnSite account, it is called the Digital ID Center. For the OnSite IPsec Account, it is called the IPsec CSR Enrollment Page.

To obtain a server certificate for the switch, you need to go to the IPsec CSR Enrollment page, paste the PKCS10 Certificate Request generated at the switch to the area provided by the CSR Enrollment Page, fill out the authentication information and challenge phrase, and then submit the request. This request needs to be approved by the IPsec Account Administrator. The administrator accesses the Control Center of the IPsec account by presenting the Administrator Certificate when connecting to the URL <https://onsite-admin.verisign.com>. When the administrator has examined and approved the request, you are sent an e-mail with the approved server certificate. You can then import the certificate into the switch by copying the base 64-encoded certificate from the e-mail to the browser's cut-and-paste buffer.

Certificate lifetime

You should issue switch certificates for a long enough time to avoid frequently reapplying for a new switch certificate. Nortel Networks recommends a validity period of at least one year or more. CA certificates generally have long lifetimes, for example, two years or more. A CA certificate must be re-imported into the switch when it reaches the end of its validity period.

Certificate use

Certificates are used each time an IPSec tunnel is established using RSA Digital Signature authentication. The client private key digitally signs the tunnel connection request. Additionally, the client transmits its corresponding X.509 certificate to the switch as part of the tunnel connection request. The switch uses a previously imported CA certificate for the CA that issued the user's certificate to verify the user's certificate.

If the certificate signature verifies, the certificate is not beyond the end of its lifetime, and the certificate has not been revoked, the public key in the user's certificate is used to verify the signature on the tunnel connection request. If this verification succeeds, the switch digitally signs a response to the connection request using the switch's private key. The switch also sends the X.509 certificate that corresponds to the private key that is used to sign the connection response. The client verifies the switch's certificate using the CA certificate in the user's Entrust certificate store. The client also checks the revocation status of the switch's certificate before using the public key in the certificate to verify the digital signature on the tunnel connection response.

For clients and switches to verify each other's certificate, each must possess a CA certificate that can be used to terminate the verification processing. The switch must have obtained a server certificate from a CA that is acceptable to the remote client. Entrust-based clients must have a copy of the CA certificate in their Entrust client certificate store generated by the initial certificate request. VeriSign clients are given applicable CA and RA certificates at installation time. Applicable CA certificates covering all CAs that can be used by remote users or the root CA in a hierarchical CA infrastructure must be imported into the switch's tunnel certificate store using the System→Certificates screen.

Certificate maintenance, expiration, and termination

Certificates expire and are not usable on the expiration date recorded in the certificate. For the switch, you can renew certificates through the same procedures used for the initial application. Renewal is automatic for Entrust and VeriSign client-based user certificates if the applicable PKI is accessible from a network, possibly using the tunnel connection itself.

End entity certificates (users or switches) can be revoked at any time by the CA. Certificate Revocation List (CRL) processing occurs automatically by the IPSec client software each time a certificate-based IPSec tunnel is established. A revoked server certificate prevents the establishment of the tunnel connection. CRL processing by the switch is optional. If in use on the switch, the revocation of a user's certificate by the CA prevents tunnel connection establishment.

Certification authority network controls

For an Entrust-based PKI, the CA must be protected behind a router firewall of sufficient strength to protect it from attacks from public or insecure networks. You can configure the firewall to allow only necessary protocols to allow the CA functions to pass through; all others are disabled.

Such access by the firewall permits Internet-based users to create an initial certificate request and response exchange without needing an initial tunnel connection to the CA. In the case of the Entrust PKI CA, TCP access to the default port 709 and the default LDAP port 389 permits clients to register and maintain their certification. Entrust uses a Secure Exchange Protocol (SEP) to protect communications between the client and CA software.

VeriSign's Onsite CA service is hosted at a VeriSign facility and is thus available over the Internet. Therefore, a VeriSign client with HTTP access to the Internet can enroll and maintain certificates.

Control of public and private keys

End entities can generate their own public key pair to allow the private half of the key pair to remain private. Private keys use a password-based method to prevent unauthorized use. The switch uses PKCS #5 and PKCS #8 storage for the private half of the key pair. It uses a password that is stored in the switch's system flash, which protects the private keys stored on the switch's system disk from unauthorized use. switch private keys can be backed up as part of the normal FTP-based backup.

You can use certificates to control IPSec tunnel connection policy. You can use CA certificates to permit authentication of all users (Allow All) who possess a certificate issued by a particular CA. Using Allow All, the switch trusts the CA to establish the true identity of a user. If the user's certificate is within the certificate validity period, and the certificate's signature can be verified using the CA certificate, and the user's certificate is not on the CA's CRL, then the tunnel connection is permitted.

You assign the user to the group associated with the particular CA. The group provides the normal connection attributes, such as hours of access and quality-of-service levels. Using the Allow All policy means that once users are certified by the CA, they can create a tunnel connection as long as their certificate is in good standing. This is analogous to using a RADIUS server to authenticate a tunnel connection.

For finer access control, you can provide individual instances for users. This means that you add the user's subject DN as it appears in his certificate to control tunnel access and group association. Although this requires more configuration, this method allows precise control over user access.



Note: If you want to switch from using a certificate with an alternate subject name to a certificate without an alternate subject name, go to the Profiles→Branch Office→Edit Connection display and from the Subject Alt Name list, select: (Do not use an Alt Subj Name). Click OK. Then, return to the Profiles→Branch Office Edit Connection display and select the certificate without the alternate subject name. Finally, select OK again.

You can also combine the two methods of access control to channel-specific individuals into specific groups, and all other individuals from the same CA into a general group. This achieves scalability with precise control. The switch always searches for a specific instance of a user before checking whether authentication permits access.

Certificate revocation list (CRL)

A CA can revoke user and server certificates whenever the associated key pair is no longer valid, the key pair has been compromised, the user has left the organization, or a server has been retired, among other reasons. When a certificate is revoked, the CA updates an associated revocation list with the revoked certificate's serial number. This list is referred to as a certificate revocation list (CRL). A CA can have one or more associated CRLs. For example, the Entrust PKI CA stores up to 750 serial numbers in one CRL before creating additional CRLs.

CRLs are published by the CA in an associated LDAP-accessible directory service. The publication frequency is set by the CA administrator. In an Entrust environment, a new CRL can be automatically published at a set time, at any time manually set by an administrator, or whenever a certificate is revoked. In a VeriSign OnSite environment, new CRLs are published at a fixed interval, typically 24 hours.

The switch can optionally use CRLs to verify the revocation status of user certificates. If enabled on the switch, CRLs are periodically retrieved from the CA's LDAP directory store and cached into the switch's associated LDAP database. This allows for rapid verification of user certificates during IPSec tunnel establishment. You can configure the frequency with which the switch checks for a new CRL.

The IPSec client uses associated CRLs to check server certificate revocation status. For an Entrust environment, the client CRL updates whenever a tunnel connection is established. In a VeriSign environment, the client updates the CRL when the existing CRL is over 24 hours old.

Because a CRL is signed using the CA's private key, it is protected against tampering. The switch verifies the CRL signature each time it is used.

CRL servers

The LDAP server that contains CRLs for the CA certificates on the switch must be reachable from the public or private interface.

The switch's LDAP database stores a copy of the CRL that the switch retrieves from the CA's LDAP server. If the switch has difficulty retrieving the CRL, an error message appears in the event log. Failure to retrieve the CRL causes the switch to retry every minute. If the switch successfully retrieves the CRL, an event log message indicates this success. If the retrieval is successful, but you still get messages indicating the CRL needs updating during user or branch authentication, there are two potential problems:

- You have an expired CRL due to lack of publication by the CA: For Entrust, go to the Entrust/Admin→CAs menu, and select Issue all CRLs. For VeriSign, contact VeriSign technical support.
- Your switch system date, time, or time zone might be incorrect, which leads the switch to believe the CRL has expired. Check the system date, time, and time zone information from the System→Date and time page.

Using digital certificates


Follow these steps to be able to authenticate remote IPSec users with digital certificates. Within each of these steps there are several substeps, all of which are detailed in a later section.

- 1 Get a Server Certificate issued by the PKI (System→Certificates: Generate Certificate Request).
- 2 Get the CA Certificate for the PKI that issues user certificates (System→Certificates: Import Tunnel Certificate).
- 3 Enable RSA Digital Signatures for the applicable remote user groups (Profiles→Groups→Edit: IPSec).
- 4 Verify that your Time Zone is correct (System→Date and Time).
- 5 Add clients into the Entrust PKI or VeriSign Onsite service
- 6 Establish an IPSec session using RSA Digital Signature authentication.

Getting a server certificate

You must obtain a Server Certificate and get the CA's certificate from the CA. You should also define the Server policy and group associations for certificate-based authentication. [Figure 49](#) shows the process for obtaining a server certificate.

Figure 49 Obtaining a server certificate



Generating a server certificate request

To generate an Entrust Server request from the Entrust Administrator:

- 1 Start the Entrust Administrator component of the Entrust PKI.
- 2 Select Add New User to preauthorize a certificate for your switch.
- 3 Supply a User Name (server name); for example, Marketing Extranet Server.
- 4 Under the Certificate General Information tab, select the category Web and the type Web Server. If the drop-down list does not show the Web selection, you have not properly installed your Web Connector license string (Users→Display/Change license information for Web).

- 5 Click Apply, then OK. A message states that you have been successful.
- 6 Click OK. From the Setup Information tab, the system provides you with the Reference number and Authorization code.
- 7 Provide this information to the switch administrator using a secure mechanism (for example, telephone or postal service).

To generate a server certificate request from a VeriSign CA Server:

- 1 Go to the System→Certificates screen.
- 2 Click Generate Certificate Request.
- 3 Enter the requested information. The first time you request a key, you are prompted for a private key password. This password is stored securely with the switch's flash memory, and it is used to protect the private keys that are stored on the switch's disk from misuse.
- 4 You must minimally fill in the Common Name, Organization Unit, and Organization fields. (You can fill in the common name with any name).
- 5 Click OK.
- 6 Select and copy the certificate request, following the displayed instructions.
- 7 Click Return.

Requesting a server certificate

To request a server certificate for the switch using Entrust:

- 1 Go to the System→Certificates screen.
- 2 Click Generate Certificate Request.
- 3 Enter the requested information. The first time you request a key, you are prompted for a private key password. This password is stored securely with the switch's flash memory, and it is used to protect from misuse the private keys that are stored on the switch's disk.
- 4 On the Create New Key and Certificate Request screen, for the Common Name enter the Entrust Reference number (for example, 00355377). If you are using Web Connector, you must also enter your organization and country. The Entrust PKI assigns the final subject name in the certificate.

- 5 Click OK.
- 6 Select and copy the certificate request, following the displayed instructions.
- 7 Click Return.

To request a server certificate for the switch using VeriSign:

- 1 Go to the IPsec CSR Enrollment Page. VeriSign provides the URL after you have a registered the IPsec account. It may be something similar to the following URL: <https://onsite.verisign.com/CustomerName/ipsecEnroll.htm>.
- 2 Provide the Certificate Signing Request (CSR), which is the PKCS10 Certificate Request generated from the switch.
- 3 Fill out the authentication information.
- 4 Choose a challenge phrase.
- 5 Submit the request. VeriSign sends the approval in an e-mail message.

Obtaining a server certificate from the CA

To obtain a server certificate from the CA using an Entrust Server (Web Connector: Non-secure Site):

- 1 Enter the Entrust Web Connector URL (supplied by the Entrust Administrator) into the address field of your Web browser.
- 2 Select Web Site Administrators.
- 3 Select Retrieve a site certificate for your Web server.
- 4 Enter the Reference number and Authorization code.
- 5 Under Options, select “displayed a PEM encoding of certificate in raw DER” (default).
- 6 Paste the encoded server certificate request into the paste box.
- 7 Click Submit Request. The Web Connector supplies you with the encoded server certificate.



Note: If you get Internal error: (2131) CMS: invalid request data, then you have entered an invalid Entrust Reference Number on the Create New Key and Certificate Request screen of the switch.

Follow the instructions on the screen to copy the encoded server certificate to your clipboard.

If you are using Entrust VPN Connector, do the following:

- 1 Copy the PKCS#10 certificate request and save it into a text file; be sure to include the Begin and End certificate request lines.
- 2 Transfer the file to the VPN Connector administrator. The VPN Connector Administrator reads the new PKCS#10 request and completes any required additional fields. When the VPN Connector administrator grants the switch's certificate request, use the screen defaults (Format:PEM and disable Wrap with PKCS-7). The VPN Connector Administrator then transmits the output file to you.

If you are using VeriSign, the IPsec Account Administrator sends e-mail to the Server Certificate Applicant (switch administrator) to indicate the approval or rejection of the certificate request. If the request is approved, the base 64-encoded certificate will be attached to the e-mail from the IPsec Account Administrator.

Importing the server certificate to the switch

To import the server certificate to the switch using Entrust:

- 1 When you receive the encoded certificate response file, open it with a text editor (such as, Notepad) and copy all of the text to the clipboard.
- 2 Click on the System→Certificates: Import Tunnel Certificate button.
- 3 Paste the encoded certificate response file into the text box.
- 4 Click OK. The server certificate entry displays on the System→Certificates: Installed Tunnel Certificate screen.

To import the server certificate to the switch using VeriSign:

- 1 Paste the encoded certificate response file into the text box.
- 2 Click OK. The server certificate entry appears on the System→Certificates: Installed Tunnel Certificate screen.

Obtaining the CA certificate from the PKI

If you are using Entrust Web Connector, enter the Entrust Web Connector URL into your Web browser's address field, select Web Site Administrators, and copy the displayed certificate to the clipboard.

If you are using Entrust VPN Connector, the CA certificate is stored in the VPN Connector installation directory (for example, Program Files\Entrust\VPN Connector) under the name vpnconccacert.pem. Using a text editor (such as Notepad), copy the certificate to the clipboard.

If you are using VeriSign OnSite, you need IPsec and Private Client OnSite Accounts from VeriSign. (Even though there are two different VeriSign accounts both accounts share the same CA). Go to the IPsec User Service Page to access the certificate. The base 64-encoded CA certificate is usually under the URL <https://onsite.verisign.com/CustomerName/ipsec>.

Importing the CA certificate to the switch

To import the CA to the switch using Entrust Web Connector, Entrust VPN Connector, or VeriSign OnSite:


- 1** Go to the Contivity VPN Switch URL.
- 2** Go to the System→Certificates screen.
- 3** Select Import Tunnel Certificate.
- 4** Select Trusted CA Certificate (default).
- 5** Paste the certificate into the paste box.
- 6** Click OK. The Installed Tunnel Certificates table displays the certificate entry.
- 7** Enable Allow All if you want to allow all users with certificates issued by this CA to authenticate with the switch, regardless of whether they have a user entry in the switch's LDAP database. By default the CA Certificate does not Allow All users authentication, so only users with their subject distinguished names (DNs) entered into the Profiles→Users screen are able to authenticate using certificates issued by this CA. If you enable Allow All users to authenticate, you must also select a group for these users from the Default Group drop-down list box. If you want only specific instances of users to

authenticate with the CA authority, you must configure each of these users from the Profiles→Users→Edit screen, and disable Allow all authentication for this CA. Only these users can then perform IPsec RSA Digital Signature Authentication using a certificate issued by this particular CA.

- 8 Click OK. You have now obtained the CA Certificate against which remote users can authenticate. Repeat this operation if multiple CAs will be issuing user certificates.
- 9 Optionally, you can configure a CRL distribution point to enable revocation checking of client certificates. Click on the System→Certificates: Installed Tunnel Certificates: CA Details button, enter the appropriate CRL Information, and click on OK.

The Enabled check box enables CRL checking of certificates for the particular CA. The Search Base, Host, Connection, and Update frequency values must be set for proper access to the CRL LDAP directory store. [Figure 49](#) shows how to enable server CRL usage.

Figure 50 Enabling server-side CRL usage



Enabling RSA digital signature IPsec authentication

You must enable RSA digital signature support for groups. Specifically, this includes any default groups that are associated with CAs, and the groups containing any specific instances of users who are doing certificate-based authentication.

- 1 From the Profiles→Groups→Edit→IPsec→Configure: RSA Digital Signature field, click on the RSA Digital Signature check box to enable.
- 2 Select the appropriate Default Server Certificate from the drop-down list box. This is the certificate that is sent to clients to authenticate the switch's identity.

This server certificate should be issued from the same Entrust PKI that has issued the remote access clients' certificates.

- 3 Click OK.

Verifying your time zone setting


On the Contivity VPN Switch:

- 1 Go to the System Date & Time screen and verify that your switch's date, time, and time zone are set correctly.
- 2 Change date, time, time zone, as appropriate.

Adding remote clients to the PKI

The client receives a certificate from the CA using the Entrust Entelligence software. If you have Internet-based clients, you must make a provision for the CA to be accessible over the Internet. The client must possess a valid Reference Number and Authorization Code for the CA for the certification request. The CA and the client attempt to make a connection, which is defined in the `entrust.ini` file on the client system.

[Figure 51](#) shows an example of certificate enrollment and maintenance.

Figure 51 Certificate enrollment and maintenance

Entrust server

Follow the Entrust instructions to add a new remote user to the Entrust PKI. Provide the Remote User with the Reference number and Authorization code via a secure mechanism (for example, postal service or telephone).

Contivity VPN Client (Entrust)

You can customize the IPSec client to allow remote users to generate new Entrust certificates through the client. To create an IPSec client installation that also installs the necessary Entrust components to do Entrust certificate-based authentication, you must include the following files in the client\custom directory:

- The Entrust DLL, which is on the Nortel Networks Contivity CD in the Client\Entrust directory, is kmpapi32.dll.
- The Entrust ini file, which was created when you created the Entrust PKI, is entrust.ini.

Contivity VPN Client (VeriSign)

You can customize the IPSec client to allow remote users to generate new VeriSign certificates through the client. To create an IPSec client installation that also installs the necessary VeriSign components to do VeriSign certificate-based authentication, you must include the following files in the client\custom directory:

- A new keyword Certinfile is defined in the setup.ini file under the [options] section to customize the Contivity VPN Client installation. This file is on the CD and is named setup.ini.
- The Cert.ini file format is used to customize the installation for the customer who uses the Contivity VPN Client in the VeriSign PKI environment. This file must be created with the name Cert.ini.

When you register for a Private OnSite Account, you can discuss the number of certificates needed. There are no specific steps to add a new remote user when the OnSite Account Administrator uses manual approval to approve certificate requests from the Contivity VPN Client one by one.

If you choose to use the passcode feature, you need to refer to the Administrator's handbook provided by VeriSign, which shows how to configure the OnSite Account to enable the passcode feature. The passcode is handed to the user in a secure fashion.

Authenticating a client with a digital certificate

The client initiates a tunnel connection and sends a signature of the IPSec request packet and the user's X.509 certificate. The server verifies the client's certificate using the CA's public key from the CA certificate stored in the server's LDAP database. If the user certificate is valid, then the signature of the IPSec packet is verified. Additionally, the server verifies the certificate against the certificate revocation list (CRL).

If the client is valid, then the server signs an IPSec response packet and sends the corresponding server certificate to the client. Once the client verifies the server's certificate using the CA's public key and checks the CRL status, the client then verifies the connection response signature using the server's public key.

If the certificate is properly verified and it is not on the CRL, the client establishes a tunneled connection.

To authenticate a client:

- 1** Launch the Contivity VPN Client using the connection wizard.
- 2** Select Authentication Type: Entrust Digital Certificate.
- 3** On the Profile selection screen, select the correct users profile file, for example, Program Files\Entrust\entrust profiles\test.epf.
- 4** Continue through the wizard to complete the connection definition.
- 5** Enter the password and click on Connect to establish an authenticated tunnel connection.

Certificate maintenance

For VeriSign client-side certificate maintenance, each time a tunnel is initiated a certificate maintenance check is performed to check the certificate's key expiration data. If the key is close to its expiration date (when 70 percent of the lifetime has passed), the Contivity VPN Client prompts the user to renew the certificate. The user can renew the certificate at that moment or delay the renewal until the next time the tunnel is initiated. VeriSign asks for renewal.

For server-side and Entrust client-side certificate maintenance, the server certificate must be replaced prior to its expiration. You can do this by deleting the certificate and requesting a new one.

Microsoft certificate services

The switch supports connections from the Microsoft Native L2TP/IPSec client in Windows 2000. In this mode, an IPSec transport mode tunnel is set up between the Windows 2000 system and the Contivity VPN Switch. Digital certificates are the primary method of authenticating the IPSec tunnel. Depending upon the network and your needs, you can configure a Microsoft Certification Authority in one of four ways:

- Enterprise Root CA is the most trusted CA in an enterprise CA hierarchy. It issues certificates to users and can certify subordinate CAs. The server must be configured with Active Directory and DNS. Certificate requests are automatically approved upon submission. Users must authenticate to the domain before processing requests on the Certificate Services Web page.
- Enterprise Subordinate CA is similar to an Enterprise Root CA, but is not the most trusted CA in a corporation. This also requires Active Directory, DNS, and parent CA. The parent may be an Enterprise Root CA, a Standalone Root CA, or an external CA.
- Standalone Root CA is also the root of a trust hierarchy. It does not require that Active Directory be configured. It can issue certificates to user, computers, and certify subordinate CAs. It often is used just to certify Subordinate CAs, which issues the user and/or computer certificates. Requests to a Standalone CA are not approved automatically. This requires an administrator. The requestor checks back later to get the certificate.

- Standalone Subordinate CA issues certificates to users and computers. Requires a parent CA

If Active Directory is installed, the Enterprise Root CA is the default option, which can be changed to one of the other three. If Active Directory is not found, then the Standalone Root CA is the default option.

To set up Microsoft certificates:

- 1** Install Certificate Services on Windows 2000.
- 2** Obtain a trusted CA certificate from the Microsoft Certification Authority* server. Select all of the text and copy it to the System→Certificates screen and click on Import Tunnel or Transport Certificate. Be sure the Trusted CA Certificate radio button is selected, paste the text into the window and click OK.
- 3** Obtain a server certificate and import it from the Microsoft Certification Authority server for the Contivity VPN Switch. Go to the System→Certificates page and click on Generate Certificate Request. Enter the appropriate information in the edit fields and click OK to generate a PKCS #7 request. This displays in a browser window for you to copy the text and paste it on the Microsoft Certificate Services screen.
- 4** On the Certificates→Certificate Details screen, configure Certificate Revocation List Checking. Enter the IP address (or host name if you have DNS configured) of the server where the CRL is stored. Make sure that the appropriate CRL boxes are checked on the Services→Available page. You can access CRLs from either the public or private interfaces. The default LDAP port is 389. Also, enter in the frequency with which you want to retrieve a copy of the CRL. Assuming that the CRL is updated by the CA immediately upon revoking a certificate, this represents the time that the switch may honor a revoked certificate. Finally, enable the Search Base to allow the switch to query the LDAP database.

Chapter 8

Configuring for Interoperability

This chapter explains the requirements and procedures for setting up different vendor hardware or software to interoperate with the Contivity VPN Switch. These instructions enable you to establish encrypted tunnels to and from the switch with the noted vendors. These requirements and procedures are subject to change based on hardware and software changes by the vendors.

Procedures are available for the following products:


- Cisco 2514 Router*, Version 11.3
- Check Point, Version 4.0 with VPN-1 and FireWall-1
- Information Resource Engineering, Inc. (IRE), SafeNet/Soft-PK Security Policy Database Editor*, Version 1.0s

Configuring the Cisco 2514 router, version 11.3

To set up the switch to establish encrypted tunnel connections as in [Figure 52](#) with the Cisco 2514 router, you should configure the Cisco 2514 as displayed with the Show Configuration command.

Figure 52 Switch and Cisco 2514 network topology

Cisco Configuration Map



The following is a show config command:

```
Cisco2514# show config
Using 1088 out of 32762 bytes
version 11.3
no service password-encryption
hostname Cisco2514
enable secret 5 $1$aSJB$Xz/o4I4IqCY.FT2RH372/1
enable password password
!
crypto isakmp policy 1
  hash md5
  authentication pre-share
  lifetime 3000
crypto isakmp key test address 8.1.10.42
!
crypto ipsec transform-set esp1 esp-des esp-md5-hmac
!
crypto map bay 11 ipsec-isakmp
  set peer 8.1.10.42
  set session-key lifetime seconds 3000
  set transform-set esp1
  match address 132
!
!
interface Ethernet0
  ip address 9.1.10.2 255.255.255.0
  no mop enabled
!
interface Ethernet1
  ip address 8.1.10.2 255.255.255.0
  no mop enabled
  crypto map bay
!
interface Serial0
  no ip address
  no ip mroute-cache
  shutdown
!
interface Serial1
  no ip address
  shutdown
!
ip classless
ip route 10.18.0.45 255.255.255.255 8.1.10.42
access-list 132 permit ip host 9.1.10.51 host 10.18.0.45
access-list 132 permit ip host 10.18.0.45 host 9.1.10.51
dialer-list 1 protocol ip permit
```

```
dialer-list 1 protocol ipx permit
snmp-server community public RO
line con 0
line aux 0
line vty 0 4
password terminal
login
end
```

Configuring the switch for Cisco interoperability

To configure the switch for Cisco* interoperability, go to the Profiles→Networks screen and click Edit.

Create any local accessible networks that you want available.

- 1** Enter the IP address for the new subnet; for example, 10.18.0.45.
- 2** Enter the subnet mask for the new network.
- 3** Click Add. The Networks Edit screen appears and shows the newly created subnet in the Current Subnets list for the named network.
- 4** Add each local subnet for which you want tunneled connections coming to or going from the switch to a Network profile.
- 5** Verify that your settings are synchronized with the Cisco (Profiles→Branch Office: Edit GROUP).


For Cisco, turn off Vendor ID and Perfect Forward Secrecy (PFS). Go to the Profiles→Groups→IPSec: Configure screen.

- 6** Create and configure the Branch Office connection on the switch, using the network profile you just created for the local accessible network. Refer to Chapter 5 for additional information on Branch Office configurations.
- 7** On the Profiles→Branch Office screen, you must enable Authentication: Text Pre-Shared Key.

Configuring the Check Point VPN-1 and FireWall-1, Version 4.0

To set up the switch to establish encrypted tunnel connections with the Check Point FireWall-1* (Version 4.0), according to [Figure 53](#), you must create certain objects on the firewall.

Figure 53 Switch and Check Point FireWall-1 network




These objects are required for the two devices to interoperate according to the preceding network configuration example:

- Go to the Manage→Network Objects→New→Network - Subnets on Check Point.
- Go to the Manage→Network Objects→New→Network - Subnet on the switch.
- Go to the Manage→Network Objects→ New→ Workstation - End point of Check Point.
- Go to the Manage→Network Objects→ New→ Workstation - End point of the switch.
- Create the Encryption rules in the firewall policy.

Establishing an encrypted tunnel on the switch side

To establish an encrypted tunnel to and from a specific network on the switch side (10.18.0.0) of the tunnel:

- 1 Select Manage→Network Objects→New: Network. The Network Properties: General dialog box opens.




- 2 Enter this general information; for example:
 - Name of the switch Subnet: PC10.18.xx.xSubnet
 - The switch's Private Subnet IP Address: 10.18.0.0
 - The switch's Private [Sub]Net Mask: 255.255.0.0
 - Select Location: Internal
 - Select Broadcast: Allowed
- 3 Click OK. The Network Objects dialog box opens.

Establishing an encrypted tunnel on the firewall side

To establish an encrypted tunnel to and from a specific network on the firewall side (9.1.10.0) of the tunnel:

- 1 Select Manage→Network Objects→New: Network. The Network Properties: General dialog box opens.




- 2 Enter the following information; for example:
 - Name of the firewall's Subnet: PC9.1.10.xx.xSubnet
 - The firewall's Private Subnet IP Address: 9.1.10.0
 - The firewall [Sub]Net Mask: 255.255.255.0
 - Select Location: Internal
 - Select Broadcast: Allowed
- 3 Click OK. The Network Objects dialog box opens.

Establishing a tunnel from the switch to the firewall

To establish an encrypted tunnel to and from the switch to the firewall:

- 1 Select Manage→Network Objects→New: Workstation. The Workstation Properties dialog box opens.




- 2 Enter this general information; for example:
 - Name of the switch Endpoint: CES
 - The switch's Public IP Address: 8.1.10.42
 - Select Location: External
 - Type: Gateway
- 3 Click OK.

The configuration of this workstation object should be done after the firewall's workstation object has been created.

Establishing a tunnel to and from the firewall

To establish an encrypted tunnel to and from the firewall to the switch:

- 1 Select Manage→Network Objects→New: Workstation. The Workstation Properties dialog box appears.



- 2 Enter this general information; for example:
 - Name of the firewall Endpoint: SQA52


- The firewall's IP Address: 8.1.10.52
- Select Location: Internal
- Select Type: Gateway
- Click on Exportable
- Click on FireWall-1 installed

3 Click OK.

Configuring the firewall's public interface

To configure the firewall's public interface:

- 1** Click on the Workstation Properties: Interfaces tab. The Workstation Properties: Interfaces dialog box opens.



- 2** Click Add.
- 3** Enter the interfaces information:
Name: Private
Net Address: 9.1.10.52


Net Mask: 255.255.255.0

- 4 Click OK. The Workstation Properties: Interfaces dialog box opens with the newly created interface listed.

Configuring the switch authentication schemes

To configure the switch's authentication schemes:

- 1 Click on the Workstation Properties: Authentication tab.




- 2 Click OS Password
- 3 Click OK.


Configuring the firewall's encryption methods

To configure the firewall's encryption methods:


- 1 Click on the Workstation Properties: Encryption tab. Here you assign the firewall responsibility for the 9.1.10.x subnet.



- 2 Select:
 - Encryption Domain: Other, and click on PC9.1.10.xSubnetwork
 - Encryption Methods Defined: click on ISAKMP/OAKLEY
- 3 Click Edit. The ISAKMP Properties: General dialog box opens.



- 4 On the ISAKMP Properties→General tab, select:
 - Encryption Method: DES
 - Authentication Method: Pre-Shared Secret
 - Hash Method: MD5
- 5 Click Edit Secrets. The Shared Secrets dialog box opens.



- 6 In the Shared Secrets List, double-click on the appropriate shared secret; for example, CES (the switch workstation object). This creates a shared secret between the switch and the firewall.
- 7 Complete Configuring the switch Workstation Object.

Configuring the switch's public interface


To configure the switch's public interface:

- 1 Select Manage→Network Objects→Select CES Workstation Object.
- 2 Click on the Workstation Properties: Interfaces tab. The Workstation Properties: Interfaces dialog box opens.
- 3 Click Add.
- 4 Enter the following information:
 - Name: CES Public
 - Net Address: 8.1.10.42
 - Net Mask: 255.255.255.0
- 5 Click OK. The Workstation Properties: Interfaces dialog box opens with the newly created interface listed.

Configuring the switch's encryption methods


To configure the switch's encryption methods:

- 1 Click on the Workstation Properties: Encryption tab. The Workstation Properties: Encryption dialog box opens.




- 2 Select the following:
 - Encryption Domain: Other, then click on PC10.18.x.xSubnet.
 - Encryption Methods Defined: Click on ISAKMP/OAKLEY, then click on Edit.

The ISAKMP Properties: General dialog box opens.




- 3 On the ISAKMP Properties→General screen, select:
 - Encryption Method: DES
 - Authentication Method: Pre-Shared Secret
 - Hash Method: MD5
- 4 Click on Edit Secrets. The Shared Secrets List dialog box opens.



- 5 In the Shared Secrets List, double-click on the appropriate shared secret; for example, SQA52 (the FireWall workstation object). This creates a shared secret between the switch and the firewall.
- 6 Enter the secret (for example, passwOrd), click set and OK. The Shared Secret List now shows asterisks for the shared secret.
- 7 Click OK. The ISAKMP Properties dialog box opens.
- 8 Click OK. The Workstation Properties dialog box opens.


Adding rules to the configurations

You must add rules to enable the configurations that you have just created. After adding the rules, the VPN-1 and FireWall-1 Security Policy screen should look like this:



- 1 Use the Edit: Add Rule menu choice. From the Rule editor, right-click on the Action column.
- 2 Select Source, and right-click on Any.
- 3 Click Add.
- 4 Select the 10.18.x.x subnet (source).

- 5 Select Destination, and right-click on Any.
- 6 Click Add.
- 7 Select the 9.1.10.x subnet (destination).
- 8 Select Edit Properties→General.
- 9 Select Action.
- 10 Click on Drop and right-click on Encryption.




- 11 Right-click and select Edit Properties.


Editing encryption for rules

Edit the encryption for each rule; for example:

- 1 If the encryption schemes are defined, click ISAKMP/OAKLEY. The ISAKMP Properties: General dialog box opens.



- 2 On the ISAKMP Properties→General screen, click:
 - Transform: Encryption + Data Integrity (ESP)
 - Encryption Algorithm: DES
 - Data Integrity: MD5
 - Allowed Peer Gateway: SQA52.
- 3 If the encryption schemes are defined, click ISAKMP/OAKLEY. The ISAKMP Properties: General dialog box opens.




4 On the ISAKMP Properties→General dialog box, click:

- Transform: Encryption + Data Integrity (ESP)
- Encryption Algorithm: DES
- Data Integrity: MD5
- Allowed Peer Gateway: CES

On the Profiles→Branch Office: Connectivity screen, the Forced Logoff attribute is equivalent to the ISAKMP SA time field on the Check Point screen (Policy→Properties→Encryption).

On the Profiles→Branch Office: Connectivity IPsec screen, the Rekey timeout is equivalent to the IPSEC SA field found on the Check Point screen (Policy→Properties→Encryption). A sample Check Point Policy→Properties→Encryption screen follows.



Configuring the switch for Check Point interoperability

To configure the switch for Check Point interoperability, go to the Profiles→Networks screen and click Edit. Create the network object used for local accessible networks; for example:

- 1 Enter the IP address for the new subnet: 10.18.0.45.
- 2 Enter the subnet mask for the new network.


- 3 Click Add. The Networks Edit dialog box opens and shows the newly created subnet in the Current Subnets list for the named network.
- 4 Add each local subnet for which you want tunneled connections coming to or going from the switch to a Network profile. Refer to Chapter 5 for additional information.
- 5 Verify that your settings are synchronized with the firewall (Profiles→Branch Office: Edit GROUP).

For some vendors, you want to turn off Vendor ID and Perfect Forward Secrecy (PFS). Go to the Profiles→Groups→IPSec: Configure screen to do this.

Create and configure the Branch Office connection on the switch, using the network profile you just created for the local accessible network. On the Profiles→Branch Office screen, you must enable the Authentication: Text Pre-Shared Key option.

Configuring the SafeNet/Soft-PK Security Policy Database Editor, version 1.0s


To set up the switch to establish encrypted tunnel connections with the IRE Soft-PK Security Policy Client as illustrated in [Figure 54](#), you should configure the screens as described on following pages.

Figure 54 Switch and IRE SafeNet network topology

Connecting to IRE SafeNet/Soft-PK Security Policy Client

To set up the switch to establish encrypted tunnel connections with the IRE SafeNet/Soft-PK Security Policy Client, follow these instructions:

- 1 Open the SafeNet/Soft-PK Security Policy Client, and click File: New. The following screen configures the network so that any packets going to the 10.18.0.0 subnet goes through the switch's 8.1.10.42 interface to establish a tunnel.



- 2 Click on the switch: CES 10.18.x.x.
- 3 Click Connection Security: Secure.
- 4 Click Remote Party Identity and Addressing:
 - ID Type: IP Subnet
 - Subnet: 10.18.0.0.
 - Mask: 255.255.0.0
 - Protocol: All
- 5 Connect using Secure Gateway Tunnel:
 - ID Type: IP Address

- 8.1.10.42


The SafeNet/Soft PX Security Policy Editor dialog box opens.

- 6 Click on My Identity to configure the SafeNet client, and enter:
 - Select Certificate: None
 - ID Type: IP Address
 - Port: All
- 7 Click Pre-Shared Key. The Pre-Shared Key dialog box appears.




- 8 In the Pre-Shared Key dialog box, click Enter Key, then enter the pre-shared key.

- 9 Click OK. The SafeNet/Soft-PK Security Policy Editor dialog box appears.



- 10 Click Security Policy: Select Phase 1 Negotiation Mode: Main Mode.

- 11 Click Enable Replay Detection.



12 On the Authentication (Phase 1), Proposal 1, Authentication screen, enable the following:

- Authentication Method: Pre-Shared key
- Encrypt Alg: DES
- Hash Alg: MD5
- SA Life: Seconds and 3000 (Seconds)
- Key Group: Diffie-Hellman Group 1

13 On the Key Exchange (Phase 2), Proposal 1 screen, enable the following:

- Encapsulation Protocol (ESP)
- Encrypt Alg: DES
- Hash Alg: MD5
- Encapsulation: Tunnel
- SA Life: Seconds and 3000 (Seconds)

Configuring the switch for IRE interoperability

To configure the switch for IRE interoperability, go to the Profiles→Networks screen and click Edit.

Create the network object used for local accessible networks:

- 1 In the Networks Edit screen, enter the IP address for the new subnet; for example, 10.18.0.45.
- 2 Enter the subnet mask for the new network: 255.255.0.0.
- 3 Click Add. The Networks Edit screen reappears and shows the newly created subnet in the Current Subnets list for the named network.
- 4 Add each local subnet for which you want tunneled connections coming to or going from the switch to a Network profile.
- 5 Verify that your settings are synchronized with the SafeNet client (Profiles→Branch Office: Edit GROUP).
- 6 Create and configure the Branch Office connection on the switch, using the network profile you just created for the local accessible network. On the Profiles→Branch Office screen, you must enable the Authentication: Text Pre-Shared Key option.
- 7 For some vendors, you want to turn off Vendor ID and/or Perfect Forward Secrecy (PFS). Go to the Profiles→Groups→IPSec: Configure screen to do this.

Third-party client installation

The switch supports third-party IPSec clients and includes support for the following:

- Authentication using either pre-shared authentication (using IKE Aggressive Mode) or digital signature certificate authentication (using IKE Main Mode) into a switch remote access user's IPSec account for third-party IPSec clients.

- Client address assignment with either the client's external IP address, or a pre-arranged internal IP address, as the address that is negotiated during the IKE Quick Mode exchange as the client's address to be used within the IPSec tunnel formed as a result of the Quick Mode negotiation.
- Split tunneling with third-party IPSec clients, such that if split tunneling is enabled on the switch, then the subnet that the client specifies as the switch's identity within the tunnel during IKE Quick Mode must be listed as one of the split tunnel networks for the Quick Mode proposal to be accepted. If split tunneling is not enabled on the switch, then the switch identity that the client specifies for Quick Mode can be any value that the client chooses.

Depending on the third-party client that you are using, you need to configure either a branch office tunnel or a user tunnel. For example, the switch has been configured and tested with the Linux FreeS/WAN* client. If you are using the FreeS/WAN Linux* client, you must configure your user and the switch as a Branch Office tunnel. If you are using another client that supports IPSec Aggressive Mode, you can configure your switch as a user tunnel.

Configuring the switch as a branch office tunnel

To configure the switch as a branch office tunnel:

- 1** Go to the Profiles→Branch Office screen and click on the Define Branch Office Connection button. The Branch Office→Define Connection screen appears:
- 2** Specify the addresses of the public interfaces of the two switches forming the connection.
 - a** For the Local Endpoint Address, click the drop-down list and select the address of the local switch.
 - b** In the Remote Endpoint Address field, enter the address of the remote switch that will form the opposite end of the branch office connection.
- 3** Set the tunnel type to IPSec.
- 4** Depending on what your third-party clients support, you can use either pre-shared key or digital certificate authentication. Click to enable the User Name and Password to authenticate user identity. The user name is the user's IP address and the password can be any password. Match the pre-shared secret with the client shared secret.


- 5 Click RSA Digital Signature to enable certificate authentication if your third-party client supports RSA Digital Signature authentication. You must then click the drop-down list box to choose a Default Server Certificate. You configure servers from the System→Certificates screen.
- 6 Go the Profiles→Branch Office screen and click on the Edit button, scroll down to the IPSec section and click on the Configure button. The Branch Office screen appears:
- 7 Select the encryption type supported by your third-party client.
- 8 Select enable or disable for the VendorID.
- 9 Set Perfect Forward Secrecy (PFS) to match the client side.
- 10 In the Rekey Timeout section, enter the amount of time to which you want to limit the lifetime of a single key used to encrypt data. The default is 08:00:00 (8 hours).
- 11 In the Rekey Data Count section, you can choose to set a Rekey Data Count depending on how much data you expect to transmit through the tunnel with a single key. The default is 0 KB; a setting of 0 disables this count.

Configuring the switch as a user tunnel

If you have third-party client software that supports aggressive mode IPSec, you can configure the switch as a user tunnel. You must use either LDAP database or certificate authentication. The switch supports both pre-shared key and RSA digital signature authentication methods and you must specify at least one of these methods.

Nortel Networks recommends enabling split tunnels for all groups that support third-party clients. If split tunneling is disabled, third-party clients can only connect if the group is configured to allow undefined networks. This means that the client can establish IPSec security associations for all networks. If you do not enable split tunneling, you must enable the Allow undefined networks option.

[Figure 55](#) shows a network with a split tunneling environment.

Figure 55 Split tunneling example

To configure the switch as a user tunnel:

- 1 Go the Profiles→Groups screen and click on the Add button. Enter a group name of up to 64 characters (spaces are permitted), for example, Research and Development.
- 2 Click on the Edit button next to the name of the new group, scroll down to the IPSec section, and click on the Configure button. The IPSec Edit screen appears.
- 3 Set Split tunneling to Enabled if you want your switch to have control over the networks that the third-party client can access. If Split Tunneling is disabled and “Allow undefined networks for non-Contivity VPN Clients” is enabled, the clients can connect to all internal networks. If you select both Split Tunneling and Allow undefined networks for non-Contivity VPN Clients, the switch uses the split tunneling feature and ignores the Allow undefined networks selection.
- 4 Under Client Selection, select “Non-Contivity VPN Clients (Linux*)” or “Both Contivity and Non-Contivity VPN Clients” from the drop-down list.
- 5 Third-party clients can use either pre-shared key or digital certificate authentication. Click the check box to enable the User Name and Password to authenticate user identity. If you are using Main Mode, the user name is the user’s IP address and the password can be any password.

Click RSA Digital Signature to enable certificate authentication if your client supports this. You must then click the drop-down list box to choose a Default Server Certificate. You configure servers from the System→Certificates screen.

- 6** Selections in the Encryption fields are dependent on the type of encryption that your third-party client supports.
- 7** Click the checkbox to enable Perfect Forward Secrecy (PFS). PFS ensures that if one key is compromised, subsequent keys are not compromised.
- 8** In the Forced Logoff edit box, specify a time after which all active users are automatically logged off. The default is 0, which means the option is turned off. The possible range is 00:00:01 to 23:59:59.
- 9** Click the checkbox to enable Compression for IPsec tunneling.
- 10** In the Rekey Timeout section, enter the time you want to limit the lifetime of a single key used to encrypt data. The default is 08:00:00 (8 hours).
- 11** In the Rekey Data Count section, you can choose to set a Rekey Data Count depending on how much data you expect to transmit through the tunnel with a single key. The default is 0 KB; a setting of 0 disables this count.
- 12** Select Enable or Disable depending on whether you want to allow IPsec Data Protection.

Chapter 9

Switch Administration

This chapter describes administration tasks that help you operate the switch. These tasks provide details on scheduling backups, upgrading the software image, saving configuration files, performing file maintenance, creating recovery diskettes, and system shutdown.

Administrators

The switch supports multiple administrators. You can assign administrative users different rights to allow them to allow or prevent them from managing or viewing switch and user configuration information. You assign administrative privileges and rights on the Profiles→User→Edit screen. The switch also supports a primary administrator.

The Administrator Settings screen allows you to change the Primary administrator user ID and password. It also controls the Administrator Idle Timeout Setting for all administrators, the default language, and serial port settings.

There can be only one primary administrator. The primary administrator user ID and password combination provides the user with this information access to all screens and control settings. The Primary Administrator User ID and Password are also used to access the serial port and the recovery disk.



Note: The primary administrator user ID and password are only saved during a system shutdown. Therefore, once you set these parameters, you must implement an Admin→Shutdown to save the new settings. Doing a reset (using the Reset button on the back of the switch) does not store the parameters.



Note: Do not lose or forget the password once the switch has been configured. Losing or forgetting your password would require you to return the switch to Nortel Networks for reconfiguration to default settings. All settings and backups would be lost. There is no way to access the system without the primary administrator password.

You can change the primary administrator user ID and password on the Admin→Administrator screen.

Backups

The switch checks at regular intervals to see whether system file changes have been made. When there are changes, they are written to each of the backup servers. All of the system files are backed up the first time; thereafter, only the files that have changed are backed up.

Backups are performed using the File Transfer Protocol (FTP). Files are transferred from the switch's hard disk to the designated path under the default FTP directory on the backup file server.

You must provide the backup Host servers' IP address, path, backup interval, user ID, Password, and Password Confirmation on the Automatic Backup screen. This information is stored in flash memory and appears on the Automatic Backup screen when you access it later. Additionally, the primary backup data is stored during the Recovery procedure (refer to the section "[Recovery](#)"), in the unlikely event that it is needed.

Some FTP servers do not use standard FTP port numbers, and they cannot be used as download FTP servers for Nortel Networks software. To enable automatic backups, go to Admin→Auto Backup screen.


If you want to save a certain configuration for a later date, you must understand that there are two components that define a given configuration: the configuration file and the LDAP database.

Saving a configuration from the Admin→Configs screen Save Current Configurations option saves only the operational parameters in the configuration file, such as interface IP addresses and subnet masks, backup host IP addresses, and DNS names.

To completely save the Contivity VPN Switch configuration on the internal LAN server, you must also save the LDAP database, which contains the group and user profiles, filters, backup file names, and so forth. Go to the Servers→LDAP screen and click Stop Server. Next, enter a file name in the Backup/Restore LDAP Database field. You should conform to the 8.3 MS-DOS naming convention and append the file name with .ldf; for example, LDAPOne.ldf.

You can use the same backup server for multiple switches. Each switch creates a unique directory based on its serial number. [Figure 56](#) shows how to properly configure your backup servers with the host name and path.

Figure 56 Backup servers 1 and 2



This table includes key information that you must enter on the Automatic Backup screen.

Host	Path	UID	Password
192.164.34.2	/backup/v101	NOCSwitch	secret
172.19.2.30	/backup	NOCSwitch	secret

Assuming the home directory for the two FTP servers is C:/software

and assuming a serial number of 01001 for the switch, then the backups would be performed to the two hosts, as follows:

```
192.164.34.2 in C:/software/backup/v101/SN01001
172.19.2.30 in C:/software/backup/SN01001
```

The switch allows you to configure regular intervals when your system files are saved to designated host backup file servers. You can designate up to three backup file servers.

Configure Automatic Backups immediately so that you do not lose system or configuration information in case of problems. You configure the Automatic Backup servers from the Admin→Automatic Backup screen.

The switch does not begin a backup for at least 5 minutes after rebooting. This time period is to allow all resources to start operating. This delay occurs even if you go into the Admin→Auto Backup screen and request that a backup be started immediately.



Note: After entering the Automatic Backup File Servers information, click on the screen and press [Alt] [Print Screen] to save the screen image to a buffer. Next, paste the image into a text file and keep it as a record of the backup file servers that you are using.

Tools

The switch supports standard IP tools such as PING, Traceroute, and ARP show and delete. You can access these tools through the Admin→Tools screen.

The PING command generates an ICMP echo-request message, which is sent by any host to test node reachability across a network. The ICMP echo-reply message indicates that the node can be successfully reached.

The Traceroute tool is used for measuring a network round-trip delay. Messages are sent per hop and the wait occurs between each message. If the address is unreachable, it uses this formula to determine how long it takes for the Traceroute to time out.

`maximum hops (30) x the wait timeout (5) x 3 seconds`

The Address Resolution Protocol (ARP) dynamically discovers the low-level physical network hardware address that corresponds to the high-level IP address for a host. ARP is limited to physical network systems that support broadcast packets that can be heard by all hosts on the network.

Recovery

The Recovery screen allows you to configure a recovery diskette for restoring the software image and file system to the hard drive of the switch in the unlikely event there is a hard disk crash. The recovery diskette is included with your switch. You can also use this screen to create additional copies of the recovery diskette, as well as to reformat a diskette.

Remove the switch's front cover to gain access to the diskette drive. Refer to the *Getting Started* for details. By booting the switch using the recovery diskette, utilities accomplish restoration of the switch's hard disk, including:

- Hard disk reformatting
- FTP access to the hard disk
- Restoration of the previously backed up software image and file system from a backup host to the hard disk
- Downloading a new, factory default software image and file system from a file server to the hard disk

These utilities are accessed via HTTP management of the switch after it has been booted off of the recovery diskette.

Using the recovery diskette

Remove the switch's front cover. Insert the recovery diskette into the drive and press the Reset button on the back of the switch. This supplies the switch with a minimal configuration utility that allows you to view the switch from a Web browser.

In the Web browser, enter the Management IP address of your switch. The Recovery Diskette screen shown in [Figure 57](#) appears, which allows you to:

- Restore the factory default configuration or the backup configuration.
- Reformat the switch's hard disk.
- Apply a new software version to the switch.
- Perform file maintenance.
- View the Event log.
- Restart the system.

Figure 57 The Recovery Diskette screen

Recovery Disk - Netscape

File Edit View Go Communicator Help

Back Forward Reload Home Search Guide Print Security Stop

Bookmarks Go to: http://

AUDIOHIG Instant Message Internet Lookup New&Cool Netcaster Welcome to Liqu Liquid Music Ne

NORTEL NETWORKS

Recovery Diskette

The Recovery Diskette allows you to reset or restore the files on your Switch. Use these features cautiously, as they delete or re-

Diskette Software Version: V02_63.14
 Diskette Software Build Date: Oct 16 2000, 22:45:16
 Hard Disk Software Version: V02_63.21
 System Serial Number: 324

Option	Action								
Restore Restore	Restore Factory Configuration ☐ Restore original factory settings. This option resets the Switch's configuration file to the original values it had when shipped from the factory. All database entries will not be altered. Important: If you choose this option, the Switch will need to be reconfigured as if it were new. Restore Backups Restore a backup image from one of the selected servers. When restoring backup files, all configuration files, internal LDAP databases, and internal files will be restored. This option should only be used to restore (or install) a complete system image to the Switch, and should not be used as a method to upgrade the system. Note: To upgrade the Contivity Extranet Switch, use the Admin->Upgrades feature of the management interface. <table border="1"> <thead> <tr> <th>Host</th> <th>Path</th> <th>User ID</th> <th>Password</th> </tr> </thead> <tbody> <tr> <td>☐</td> <td> </td> <td> </td> <td> </td> </tr> </tbody> </table>	Host	Path	User ID	Password	☐			
Host	Path	User ID	Password						
☐									
Reformat hard disk Reformat	Formats the hard disk in the Switch. Use this option cautiously. It will destroy all the information on the Switch's hard disk.								
Apply new version Apply	Changes the version of software executing on the Switch. Use this option to change to other software versions which exist on the Switch's hard disk. When applying a new software version, the current version will be preserved under the previous version. Select the desired software version: (No version selected)								
Perform file maintenance Files	Presents a listing of directories and files on the Switch.								
View event log View	The Event log allows you to see system Events that have occurred on the Switch. This log should be used to resolve problems that occur on the Switch.								
Restart system Restart	To restart the system, remove the diskette and press the Reset button on the back of the Switch.								

Document: Done

To use the recovery diskette:

1 Restore the configuration:

To restore the factory default configuration or the backup configuration, select the hard disk drive to which you want to restore the system files; either ide0 (drive 0) or ide1 (drive 1) and then do one of the following:

- Restore the factory configuration by selecting Restore Factory Configuration, then click Restore to return the switch to its original factory default configuration. This erases data contained in flash memory and also in the configuration file.



Warning: Selecting this option requires you to rebuild your entire switch configuration again from scratch

An online message specifies the result of the Factory Configuration reset action.

- Or you can restore the switch's previously backed up configuration by clicking Restore. If you previously chose to automatically backup the file systems, then the Backup Server Host (or IP address) and Path Name, User ID, and Password appear in the table.

Click the radio button of the preferred backup server. The backed up file system, including software image and configuration files, from the latest backup copy residing on the designated server is restored onto the hard drive of your switch.

You can use the same backup server for multiple switches. Each switch creates a unique directory based on its serial number. The following example shows the Host, Path, and Serial Number (where the serial number [SN] is five digits):

```
C:/software/backup/v101/SN01001
```

The Serial Number is used to differentiate backup configurations from multiple switches that are saved on the same backup server. The Serial Number uniquely identifies each switch's backup data.

A blank row in the server backup field always appears to allow you to manually enter a backup server in case you did not configure automatic backup server locations.



Note: Because FTP servers are often different, there may be some information in your server documentation about setting paths that can help you with the upgrade procedure.

Alternatively, a new factory default software image and file system can be restored to the switch's hard disk. Specify the name or address and path of the network file server onto which the software from the Nortel Networks CD has been installed.



Note: This restores the disk to an operable but “clean” condition (for example, configuration values are at factory defaults).

To view your switch's Serial Number when the switch is operational, choose Status→System from the Navigational menu. The Serial Number is also on the bar code label on the back of the switch.

- 2 Click Reformat your switch's Hard Disk if you must reformat the hard disk, for example if you:
 - Have problems restoring your configuration that are not caused by the network or the file/backup server from which the file restoration is being retrieved
 - Want to reconfigure the switch from scratch
 - Install a new disk



Caution: Selecting this option completely wipes out anything that was stored on the hard disk.

An online message indicates whether the Reformatting of the Hard Disk was successful.

- 3 Click the list to view the available software image and file systems that are stored on the hard disk and select the image version that you want to activate.

- 4 Click Apply to apply the new version and reboot automatically. Changes are active. The switch boots to that version until changed.
- 5 Click Files to bring up the File Maintenance screen, which allows you to view the entire hard disk file system.
- 6 Click View to display the Event Log beneath the Recovery Diskette screen. This is especially useful if a Restore operation fails.
- 7 Set the boot disk by clicking the list to select the hard disk drive from which you want to boot the switch; either ide0 (drive 0) or ide1 (drive 1). Click Set.
- 8 Click Synchronize to immediately synchronize the primary and secondary disks. Thereafter, the disks automatically synchronize every hour. The switch does not synchronize the software and configuration. Everything under the system directory is synchronized except for the core directory. Synchronization happens automatically so you do not have to initiate it.
- 9 Upgrade the system boot software by clicking the list to select a drive onto which you want to update the system boot software. Click Upgrade to rewrite the boot software onto the hard disk. You would do this if the system boot sector were to become corrupted.
- 10 Restart the system by removing the diskette and pressing the Reset button on the back of the switch. Then reposition your Web browser to the Management IP address, and choose Reload or Refresh from your browser menu to access the management page of the software running on the hard disk.



Note: This procedure cannot be used for the Contivity VPN Switch 1000 due to the lack of a floppy drive in the unit.

Upgrades

You can download the latest Nortel Networks software for the Contivity VPN Switch using File Transfer Protocol (FTP). In addition to retrieving software, you can select which version of software to run.

The Host, Path, User ID, and Password fields are already filled in. You only have to select the version that you want to download. Click the View link to go to the Nortel Networks Web site where you can select the latest software version.

Additionally, you can enter the necessary information to download software from your own FTP site. This situation assumes that you have previously downloaded the software from Nortel Networks to your organization's FTP site (or loaded it from the Nortel Networks CD). In this case, simply write over the listed Nortel Networks details.

A maximum of four versions are supported on the Contivity VPN Switch system disk. If four versions exist on the Admin→Upgrade screen, you must delete one using the Admin→File System before you FTP another release.

You cannot upgrade switch software through a branch office tunnel that is translating the management address with dynamic NAT.



Note: Some FTP servers do not use standard FTP port numbers, and they cannot be used as download FTP servers for Nortel Networks software. For more information, contact Nortel Networks Customer support.

To initiate an upgrade, go to the Admin→Upgrade screen. After the new version has loaded, select “Click here to complete the upgrade” link. If the link reappears off of the main welcome page, it means there is a browser caching problem and you should reload the page.

System configuration

You can save the current or delete existing system configuration files through the Admin→Config screen. Additionally, you can select one of the previously named configurations and restore it as the current configuration.

File management

You can navigate through the switch file system through the Admin→File System→File System Maintenance screen. It lists the devices (drives) and directories. This provides flexibility in viewing details of a file or directory, and it allows you to delete unnecessary files. For example, if you had problems

performing an FTP transfer with a specific file, you could view the file details to learn its file size and when it was last modified for troubleshooting purposes. Additionally, you can toggle between hard drives when a backup drive is available.

SNMP

The SNMP screen (Admin→SNMP) allows you to configure the switch to generate SNMP Version 1 Traps, based on MIB II. You can do the following:

- Designate the remote SNMP management stations that are authorized to send SNMP Gets to the switch.
- Designate the trap hosts to which the traps can be sent.
- Configure the traps.

The SNMP counters measure packet attributes that are based on the outer IP header. In the tunneled environment there is also an inner IP header, but this IP header does not contribute to the SNMP MIB counters. For example, the outer packet header might be a good packet header and counted, but the inner packet header might be corrupted and would not contribute to the drop counter.

You can view the Health Check screen for the results of SNMP Traps.



Note: A Nortel Networks proprietary MIB is included on the Nortel Networks CD. Click on the CesTraps.mib file to load the MIB. See [Appendix A, “MIB Support”](#) for a description of CesTraps.mib.

Shutdown

The Shutdown options allow you to shut down immediately, to wait until current users are logged off, or to wait until a designated time. A normal shutdown safely terminates connections so that no data is lost, compared with a spontaneous loss of power.

Additionally, you can select whether to power off or restart after shutdown, and also choose the configuration file to use upon restarting. To allow you to conduct an orderly shutdown, you can disable new logins, and you can disable logins after the shutdown to perform system maintenance.

Always use the System Shutdown screen to shut down the system rather than the Power or Reset buttons on the back of the switch. This ensures the integrity of your file system.



Note: After performing a system shutdown, click the Reload/Refresh button to see the latest switch information.

Go to the Admin→Shutdown screen to turn off the switch.

You can either boot in normal mode or safe mode. Each mode has its own system image and configuration information. A system booted in safe mode can only accept secured management tunnel connections. Once this tunnel is established Telnet, HTTP, and FTP traffic is allowed, but no other VPN traffic is allowed through the secured management tunnel or the switch. In normal mode, the system operates with the normal software and configuration and transports both VPN traffic and management traffic.

The system can be booted in safe mode when it is requested at the system shut down or when the system is not shut down normally. Once in Safe mode, the system stays in that mode for a configurable amount of time (default is 5 minutes). If there is no administrative user logged into the switch over a secured management connection or the serial console within this time period, the system reverts back to Normal mode.

You can configure Safe mode settings on the System→Settings screen.

Switch status

The System Status screens allow you to see from the Web interface which users are logged on, their traffic demands, and a summary of your switch's hardware configuration, including available memory and disk space.

The status screens include:

- Sessions
- Reports
- System Status
- Health check
- Statistics
- Firewall
- Accounting

The switch has several logs that provide different levels of information, including:


- Event log
- System log
- Security log
- Config log

The logs are stored in text files on disk and they indicate what happened, when, and to which user (IP address and user ID).

The Event log captures real-time logging over a relatively short period of time (for example, the Event log could wrap its 2000 possible entries in minutes). The System log captures data over a longer period of time, up to 61 days.

Most events are sent to the Event log first. Significant events from the Event log are sent to the System log. Not all data that is saved by the System log comes from the Event log, but that is generally the case. The switch filters from the System log security entries for the Security log and configuration entries for the Configuration log.

[Figure 58](#) shows a Nortel Networks logging scheme.

Figure 58 Nortel Networks logging scheme

The different Log options allow you to write specific event levels to the log files and view them, including:

- Normal
- Urgent
- Detailed
- All

Sessions

You can monitor which users are tunneled into the switch, when they logged in, and the number of bytes and packets they have transmitted or received. Additionally, you can see selected session details, and you can even log off users.

Reports

The Reports feature allows you to generate comprehensive reports of users and related information. You generate reports in an on screen tabular format, and you can import them into a spreadsheet or database through the comma-delimited format.

At midnight (12:00 a.m.), the data collection task performs summary calculations and rewrites history files, along with other management and cleanup functions. You should leave the switch running overnight to perform this task. The switch must be running at midnight to generate a historical graph for the day.

The Status Reports screen allows you to view system and performance data in text or graphical format. You can generate current or historical graphs of valuable system data. The Reports feature provides a comprehensive display or down-loadable reports on user activity.

If you have multiple switches throughout the world, you might want to use the Greenwich Mean Time (GMT) standard. This synchronizes the various log files so that the timestamps are directly comparable.

System

The System Status screen shows the switch's up time, software and hardware configurations, and the current status of key devices. When there is a pending shutdown or an IPX Public Network Address change that requires a reboot, such events are listed at the top of this screen.

Health check

The Health Check screen provides an overall summary of the current state of the switch's hardware and software components at a glance.

Statistics

The Statistics screen provides many subscreens with a wealth of general and diagnostic information about the switch hardware, software, and connections. Much of the information is specifically designed for Nortel Networks Customer Support personnel to assist them in diagnosing problems. Some screens, however, such as the LAN Counters, Interfaces, and the WAN Status, might provide you with some traffic information.

Firewall

The Firewall screen shows the details of the firewall monitoring session. It is only used for the optional integrated Check Point firewall. This screen provides the same information as the following Check Point firewall command:

```
FW stat -l
```

The FW-1 not loaded message shows if you have specified either the Contivity Stateful Firewall or No Firewall on the Services→Firewall screen.

Accounting

The accounting log provides information about user sessions. The log provides last and first names, user ID, tunnel type, session start and end dates, and the number of packets and bytes transferred. You can search the log according to most of these fields.

Event log

The Event log is a detailed recording of all events that take place on the system. These entries are not necessarily written to disk, as with the System log. The Event log retains all system activity in-memory but only the significant entries are saved in the System log (on disk).

The Event log includes information on tunneling, security, backups, debugging, hardware, security, daemon processes, software drivers, interface card driver events, and so forth.

As the Event log adds in-memory information, its oldest entries are overwritten. The Event log retains the latest 2000 entries, and discards old entries when it is refreshed.

System log

The System log contains all system events that are considered significant enough to be written to disk, including those displayed in the Configuration and Security logs. Events that would appear in the System Log include:

- LDAP activity
- Configuration activity
- Server authentication and authorization requests

Security log

The Security log records all activity about system or user security. The Security log lists all security events, both failures and successes. The events can include:

- Authentication and authorization
- Tunnel or administration requests
- Encryption, authentication, or compression
- Hours of access
- Number of session violations
- Communications with servers
- LDAP
- RADIUS

Configuration log

The Configuration log records all configuration changes. For example, it tracks adding, modifying, or deleting configuration parameters:

- Group or user profiles
- LAN or WAN interfaces
- Filters

- System access hours
- Shutdown or startup policies
- File maintenance or backup policies

Chapter 10

Troubleshooting

This chapter introduces the concepts and practices of advanced network configuration and troubleshooting for the Contivity VPN Switch. Its purpose is two-fold: to provide configuration details for you to consult when setting up or modifying your extranet, and to serve as a resource when diagnosing client and network problems.

Typically, there are three types of problems to address when managing an extranet:

- Connectivity
- Performance
- General

As a network administrator, your primary concern is maintaining connectivity. For extranet access, this means maintaining the secure connections between your remote users and the private intranet serviced by the switch. Performance is another area of concern. Paying attention to performance can help you address issues before they become problems.

Connectivity problems occur when the remote user cannot establish a connection with areas of their private corporate network. Because extranet connections are typically made over the Public Data Network (PDN), there are several points of failure to consider when diagnosing connectivity problems. Problems can range from something as simple as a modem configuration error on the client workstation to a complex HDLC protocol error on the T-1 WAN interface.

Troubleshooting remote access problems typically starts at the client end when the remote user cannot establish a connection, loses a connection, or has difficulty browsing the network or printing. When connectivity problems occur and the source of the problem is unknown, it is usually best to generally follow the OSI network architecture layers. Therefore, start diagnosing the physical environment, the modem, and cables before moving up to the network and application layers, (for example, pinging a host and Web browsing).

As with connectivity, there are many places in the extranet network where network performance can be impacted. By regularly checking your network statistics, logs, and health check information, and by informing users of good network practices, you can often avoid problems and enhance the productivity of your extranet.

General problems are categorized here as problems other than those related to connectivity or network performance. For the latest release-specific problems, check the release notes.

Troubleshooting tools

For the Contivity VPN Switch administrator, a robust troubleshooting “toolbox” is filled with both standard and special tools for diagnosing network problems. Standard tools like Telnet, PING, Trace Route (tracert.exe), sniffers, and analyzers are a basic necessity. To this collection, some special tools are added to the Contivity VPN Switch manager and remote client applications. These special tools include client- and switch-based tools.

Client-based tools

IPSec Contivity VPN Client Monitor provides network statistics on device, connection, and network errors that are helpful for monitoring traffic flow and assessing IPSec connection performance. Statistic counters are updated once a second. For more information on the IPSec Contivity VPN Client Monitor, refer to the Contivity VPN Client online Help.

Microsoft PPTP Dial-Up Networking Monitor provides network statistics on device, connection, and network protocols that are helpful for monitoring traffic flow and assessing PPTP connection performance. For more information on the PPTP Dial-Up Networking Monitor, refer to the PPTP Help or your Microsoft PPTP client documentation.

Switch-based tools

Manager Status→Health Check screen provides a detailed picture of how the switch is performing. View colored status indicators to evaluate individual component status, and click on associated hyperlinks to jump to manager screens for corrective action.

Manager Status→Statistics screen allows you to delve into the inner workings of the switch where you can view detailed system and network statistics. For more information on the Statistics screen, refer to “Administration.”

Manager Status→Security, Config, System, and Event Log screens allow you to view various logs recording system and network events that help you trace problems and determine their origins. For more information on the Logs screens, refer to *Reference for the Contivity VPN Switch*.

SNMP Traps allows the scripting of SNMP alerts so that a combination of system variables can signal an SNMP trap. When a trap occurs, the Nortel Networks icon on management stations turns red, signaling a problem. For more information on SNMP traps, refer to *Reference for the Contivity VPN Switch*.

Other tools

Microsoft Windows 95/Windows 98/NT Client Troubleshooting Tools--[Table 18](#) lists the tools that are helpful for diagnosing connectivity problems from Windows 95, Windows 98, and NT workstations:

Table 18 Troubleshooting tools

Windows 95/Windows 98	Windows NT	Use for...
Winipcfg command	Ipconfig command	Obtaining IP address, DNS, WINS information
Netstats command	Netstats command	Viewing statistics from Microsoft TCP/IP stack
Ping and tracert commands	Ping and tracert commands	Testing connectivity, name resolution, route tracing
Dial-Up Monitor status	Dial-Up Monitor status	Viewing modem settings, throughput and errors

The toolbox also includes the NetMedic* application from VitalSigns Software. This tool is helpful for route tracing and monitoring performance over a Web browsing connection.

Solving connectivity problems

This section lists many of the common connectivity problems that can occur and the recommended solutions. Problems, and some typical client user responses that can help with diagnosis, are categorized as follows:

Modem and dial-up problems

“I cannot browse the Web or check my e-mail over my dial-up connection.”

“I cannot ping my ISP site.”

Extranet connection problems

“I can browse the Web over my dial-up connection, but I cannot log in to my network over the extranet connection.”

Problems with name resolution using DNS services

“I logged in to my corporate network, but I get messages saying the host is unknown.”

“I can ping the host using its IP address, but not using its host name.”

Network browsing problems

“I cannot browse the corporate network.”

“I cannot print.”

“I cannot access the Internet over my extranet connection?”

Diagnosing client connectivity problems

A connection can fail at varying points in an extranet. If remote users have a problem accessing their corporate network, and the source of the problem is unknown, Nortel Networks recommends that they follow these steps to first determine whether the problem is with their modem, PPP dial-up, or with the extranet connection:

- 1** Confirm that the remote user’s modem is attached and working properly by having them run a terminal emulation program at their remote workstation, such as, Hyperterminal*, and issuing the AT command. The response back should be AT OK if the modem is operating correctly.
- 2** Verify that the remote user has a PPP dial-up connection over the internet. To do this, before they try to establish an extranet access or PPTP connection, have them try Web browsing www.nortelnetworks.com or another Web site that they commonly access. If the remote user can access the Web site, their PPP dial-up connection is working properly. Refer to the section [“Common client connectivity problems”](#) to further troubleshoot the connection problem. If the remote user still cannot verify that their dial-up connection is working properly, continue with step 3.
- 3** Ask the remote user to check that their modem type and settings are configured properly. To do this, they need to right-click on the Dial-Up Networking connection icon (the icon they click to dial their connection) on

their desktop to view its properties. Verify that these settings are correct for their modem configuration.

- 4 If the remote user is connected, but unable to access any resources or servers, have them check their system's connection information by going to the Start menu, selecting Run, and typing winipcfg in the text box (or ipconfig if using Windows NT). Ask them to view the statistics for their PPP adapter and confirm that the entries match those provided by their ISP.
- 5 If the remote user is still unable to view resources or servers over their PPP dial-up connection, contact their ISP to see if they have logged any connection attempts from the user, and for additional troubleshooting assistance.

Common client connectivity problems

Extranet connection problems

If the client is successfully connecting to their ISP, but is having problems accessing their intranet over their PPTP or IPsec Contivity VPN Client connection, have them check the following areas to further troubleshoot their connection problem.

The following messages and their associated cause and action statements are directed to the IPsec Contivity VPN Client user at the remote workstation. This information is also available in the Contivity VPN Client online Help.

Remote host not responding

Cause: This indicates that the switch never responded to the IPsec connection attempt.

Action: Verify that the switch is accessible by pinging the host name or IP address that you filled into the destination field. To PING a host called extranet.corp.com, for example, open an MS-DOS* command prompt and type ping extranet.corp.com. If you receive a reply message, it indicates that the switch is accessible but it is not responding. If you received a message that says "Request Timed Out" from the ping command, it means that the switch is inaccessible. You may be able to further diagnose the problem using the MS-DOS Trace Route command (tracert.exe) on Windows 95, Windows 98, and Windows NT systems.

The switch only allows a certain number of PING packets from another Internet host before requiring a tunnel connection to be established.

Maximum number of sessions reached

Cause: This indicates that the maximum number of users for the account you are using are currently logged in.

Action: If you are the only user with access to your account, it is possible to get this error if you restarted an IPSec connection immediately after losing the dial-up connection to your ISP. This is because the switch takes up to one minute to determine that your connection has been dropped and log you off from your account. Simply wait a minute retry your connection.

Login not allowed at this time

Cause: This indicates that your account has been limited to specific hours of access and you are trying to connect outside of the allowed time.

Action: Contact your network administrator if you are unsure of your specific hours of access.

Authentication failed

Cause: The IPSec Username is incorrect or the Password is invalid for the Username entered.

Action: Verify that the user name you entered is correct and retype the password before trying the connection again.

No proposal chosen

Cause: The switch you are connecting to is not configured to handle the Authentication method configured under the current Connection Profile.

Action: Verify that you are using the correct Authentication method. If your network administrator gave you only a username and password to log in to the switch, check that Username and Password is selected under the Authentication Options, which is located under the Options menu. If you are using a Token Card or the Password Authentication Protocol (PAP) for RADIUS support, verify that you have Use Group Security Settings selected under the Authentication Options and that the correct Group Authentication Option is selected.

Other IPSec errors

Cause: Typically other error messages indicate an error in configuration on the switch that must be corrected by the network administrator.

Action: Contact your Network Administrator with the specific error message.

Extranet connection lost

If the PPTP or IPSec Contivity VPN Client connection was initially established and then fails, one of two error messages appear: “The physical connection has been lost” or “The secure extranet connection has been lost.”

The physical connection has been lost

Cause: The PPP connection to your ISP was disconnected.

Action: Re-establish the PPP dial-up connection to your ISP before you re-establish the extranet connection to the remote network.

The secure extranet connection has been lost

Cause: For IPSec only, the switch that you are connected to has either logged your connection off or the switch is no longer responding.

Action: Re-establish the extranet connection by clicking the Connect button. If this works, the connection was probably lost due to the Idle Timeout configured on the switch. If no data is transferred through the extranet connection for a long period of time, normally 15 minutes or more, the switch automatically disconnects the connection.

If you were unable to successfully re-establish the extranet connection, the dial-up connection may be preventing data from traveling between the Extranet VPN Client and the switch. Hang up the dial-up connection and reconnect before you try to re-establish a connection to the switch. If you are still unable to connect to the switch, open an MS-DOS Command Prompt and try pinging the switch using the host name or address that you specified in the Destination field. If you receive a Destination Unreachable error message there is a routing problem at the ISP. If you receive a Request Timed Out error message, the switch is probably not available, and you should contact your network administrator.

Auto disconnect closes the dial-up connection during data transfer activity

Cause: In Windows 95 only, Microsoft's Auto Disconnect feature does not recognize data activity unless it passes through Internet Explorer. Microsoft has documented this as a known problem in Windows 95.

Action: At the remote workstation, disable the Auto Disconnect feature if you are not using Internet Explorer to access data on the remote network. To do this, open the Control Panel and choose the icon labeled Internet. Select the Connection property sheet and deselect the "Disconnect if idle for" box.

Problems with name resolution using DNS services

When the client can ping a host using an IP address, but not with its host name, or receives messages that the host name cannot be resolved, DNS misconfiguration is usually the problem.

Cause: A DNS server may not be configured for PPTP or IPSec connections on the switch.

Action: Validate that the Contivity VPN Client has been configured with a DNS entry. For Windows NT 4.0, open a Command Prompt and enter `ipconfig/all`. Verify that a DNS server entry is listed. For Windows 95, from the Start menu on the task bar, select Run and enter `winipcfg`. Select Nortel Networks Extranet Switch Extranet Access Adapter from the list of adapters and click More Info. Record the information displayed under DNS Server entry and verify it with the network administrator.

Cause: The hostname being resolved has both a public and a private IP address. This is commonly referred to as a split-horizon DNS.

Action: Open a command prompt and ping the host you are trying to reach with a fully qualified host name (for example, `www.nortelnetworks.com`). If you receive a response, verify that the IP address returned on the first line (for example, `www.nortelnetworks.com [207.87.31.127]`) is an IP address from the remote corporate network. If it is not, notify your network administrator that the internal hostname should be modified so that it is not the same as the external hostname.

Cause: The retail release of Windows 95 contained a bug that prevented use of more than one DNS server. This problem was fixed in OS Release 2.

Action: If you are using a release earlier than OS Release 2 of Windows 95, a patch is available from Microsoft to upgrade the `winsock.dll`. This patch is downloadable from www.microsoft.com.

Network browsing problems

Cannot browse the network (with NetBEUI)

Cause: For both PPTP and IPSec, the switch does not currently support the NetBEUI protocol.

Action: To be able to browse resources on a remote domain through a connection to a switch, it is necessary to remove the NetBEUI protocol. By removing NetBEUI, the Microsoft Client uses NetBIOS over TCP/IP to browse network resources. This applies to both the PPTP dial-up client provided by Microsoft and the Contivity VPN Client provided by Nortel Networks.

Cannot access Web servers on the Internet after establishing an Contivity VPN Client connection

Cause: For both PPTP and IPSec, this condition occurs as a result of having all network traffic passed through the corporate network. Typically, firewalls and other security measures on the corporate network limit your access to the Internet.

Action: The switch administrator can set up a default route on the switch to forward traffic to the Internet. If this default route is not configured, you need to disconnect the extranet connection to Web browse the Internet through your ISP connection.

Alternatively, if you are using a proxy-based firewall, you must set your Web browser to use the firewall to proxy for HTTP traffic when your tunnel connection is in use.

Cannot access network shares after establishing an extranet access connection

Cause: PPTP and IPsec--A Windows Internet Name Service (WINS) Server may not be configured for PPTP or IPsec connections on the switch.

Action: Validate that the Contivity VPN Client has been configured with a WINS server. Follow the steps outlined above under [“Problems with name resolution using DNS services”](#) to run ipconfig at a command prompt on Windows NT 4.0 and to run winipcfg on Windows 95. Verify that a primary WINS server is listed under the section for the adapter named IPsecShm on Windows NT 4.0 and on Windows 95 verify that a primary WINS server is listed in winipcfg for the Contivity VPN Client adapter. If there is no Primary WINS server listed, notify the network administrator that the switch may not be properly configured.

Cause: Your system may be set up for a different domain other than the one on the remote network.

Action: Skip the initial domain login when Windows 95 starts and choose Log on to the Remote Domain under the Options menu of the Contivity VPN Client dialog box. You are then prompted to log into the domain of the remote network after the extranet connection is made. This is the recommended method for users with docking station configurations.

Alternatively, on NT 4.0, Windows 98, and Windows 95, change your workstation to be a member of a Workgroup instead of a Domain:

- 1 From the Start Menu, select Settings→Control Panel. In the Control Panel, double-click on the Icon labeled Network. The Network Control Panel applet opens.

- 2 Select the Identification tab. In Windows 95, the entries on the Identification tab can be modified directly, while on NT 4.0 it is necessary to click the Change button to change the entries.
- 3 Change to use a Workgroup and verify that the computer name does not match the entry on the remote network. The name for the workgroup is not important, so enter anything.
- 4 Click OK to save the changes and reboot the machine.
- 5 When accessing a resource on the remote domain, if you are prompted for a user name and password, you need to have a domain name precede your User ID. For example, if your User ID is JSmith and you are accessing a machine on the remote domain named CORP, you would enter your user name as CORP\JSmith.

Diagnosing WAN link problems

WAN link problems can occur between the switch and the Public Data Network (PDN) at three levels:

- 1 T-1/V.35 Interface
- 2 HDLC framing
- 3 PPP Layer

If a connectivity problem occurs with the WAN link, there are two approaches you can use to diagnose and correct the problem.

- Start from the bottom—Verify that physical connectivity exists, then make sure that the HDLC link is up, and finally examine the PPP status to see if it is passing IP packets back and forth.
- Start from the top down—Go in the opposite direction, looking at PPP first and working down to the physical connection. An important point to remember when taking this approach is that at the higher protocol layers, there are more options to misconfigure, but changing them is easier and generally involves less effort.

A key point to remember when diagnosing WAN link problems is that the T-1 service provider should usually be involved in the troubleshooting effort. This is not only because they can help diagnose the problem, but also because an ISP can bring down a link if it detects errors on the line. You should notify the ISP administrator if you are planning on working on the link.

Check the T1/V.35 interface

To diagnose a problem at the WAN physical layer, use the following steps to verify that the T-1/V.35 interface to the Public Data Network (PDN) is operating correctly, and that the T-1 line is properly connected:

- 1** Have your ISP run a loopback test from their end to the CSU/DSU to verify that the external line is working correctly.
- 2** Check the connections between the switch and the CSU/DSU. The V.35 cable should be a straight-through cable and firmly seated. Also the CSU/DSU should be configured to use internal clocking, NRZ encoding with CCITT CRC for the checksum.
- 3** Make sure that all the control signals are asserted (CTS, DCD, DSR, RTS, and DTR). You can check these signals on the switch from the Manager WAN Statistics screen. If any of these signals are incorrect, you can try disabling/enabling the link from the Manager WAN Interfaces screen, or unplugging and plugging in the link. If these steps do not resolve the problem, try switching ports on the same card, switching cables, or switching to a new card, if available.
- 4** If the previous steps fail to resolve the problem, and you still suspect a problem with the physical connection, try rebooting the switch to reinitialize the WAN interface.

Check the HDLC framing

Assuming that the T-1/V.35 interface is operating correctly, use the following steps to determine whether the HDLC layer is up and running properly, and to provide information for Nortel Networks Customer Support for further diagnosis:

- 1** Check to make sure that there are no input or output errors reported on the Manager WAN statistics screen. Also look to see if the input and output counters are incrementing at all. If the input/output counters are not incrementing, or are incrementing by huge amounts, then there are probably

framing or timing errors on the link. Also, a large percentage of input errors may indicate a problem with the FCS (Frame Check Sequence) calculation.

- 2 Examine the Manager Statistics event log with debugging enabled. Any WAN-related log messages probably indicate some sort of error.
- 3 Report any of the preceding errors and messages to Nortel Networks Customer Support for assistance in diagnosing the HDLC framing problem.

Check the PPP layer

If the WAN link appears to be passing frames back and forth, yet IP packets are not flowing, the problem may be with how PPP is configured.

To examine the state of the PPP connection, and to provide information for Nortel Networks Customer Support for further diagnosis:

- 1 Check to see whether the state of the PPP connection is changing at all by periodically clicking the Refresh button while viewing the WAN statistics screen. If the state is always Down, PPP might not know that the link is up. If the state toggles between Dead and LCP Negotiating, PPP is trying to come up but cannot. This is probably due to a problem with the underlying layers, although it could also be a bad configuration of the LCP options.
- 2 If the connection fails during Authentication, then try disabling the PPP Authentication settings. A problem during Network Negotiating is usually due to misconfigured IPCP options.
- 3 Verify that all the Authentication settings match the ISP-recommended router configuration.
- 4 If the PPP layer still does not come up, enable the interface debugger to generate large amounts of packet traces in the event log. Report this information to Nortel Networks Customer Support for further diagnosis.

Hardware encryption accelerator connectivity

If the hardware encryption accelerator fails, all sessions are automatically moved over to be handled by the software.

Solving performance problems

This section describes ways for improving the performance of the remote workstation connection to the corporate network through a switch. It also includes Microsoft networking and client setup and operation tips.

Eliminating modem errors

Modem hardware errors can impact performance when connecting to your corporate network over a dial-up connection. If modem hardware errors are occurring, try the following techniques to correct these errors and improve performance:

- Adjust the modem speed—If the speed of the modem is set too high it can cause hardware overruns. Reset the modem speed to match the real speed of the modem.
- Disable hardware compression—The data passed through the extranet connection is encrypted, and encrypted data is typically not compressible. Depending on the algorithm the modem is using to compress the encrypted (non-compressible) data, the data may expand in size and overrun the modem's buffers.

Using DES encryption instead of Triple DES when using IPSec

The type of encryption used to encrypt the data passed through an extranet connection can affect performance. For IPSec, Triple DES is a stronger encryption method than DES and so the computational requirements of encrypting data with Triple DES are higher. Only the network administrator can make this change.

Hardware encryption accelerator performance

Expanded bytes and packets counts reflect packets that expanded when compressed by the LZS algorithm. These counters are important because they indicate a heavier load on the accelerator since packets that expand must be sent a second time to the accelerator with compression disabled. If there are many sessions transporting incompressible traffic (such as a video stream), the overall performance of the switch degrades relative to its performance when all sessions carry compressible data (such as FTP or text files).

Performance tips for configuring Microsoft networking

For Microsoft networking to work as designed over the extranet, each of the following components, if configured, must be working together:

- DHCP Server—assigns IP addresses to clients.
- WINS Server—provides a translation of the NetBIOS domain name to the IP address.
- DNS Server—provides a translation of the IP Host name to the IP address.
- Master Browser—an elected host that maintains lists of all NetBIOS resources.
- Domain Controller, which maintains a list of all clients in the NetBIOS domain and manages administrative requests such as logins.
- Contivity VPN Switch—terminates tunnels and routes Microsoft networking requests.

The following questions and answers are particularly directed toward the WINS server and browsing issues that can help you verify whether you have correctly set up these components.

What needs to be configured on the switch for network browsing?

In the groups profiles, set the values of the DNS server and the WINS server. Remember that these are inherited values so that if all subgroups of a given group use the same servers, it is sufficient to configure them in the parent group.

If these servers are not on a directly reachable subnet from the switch, or accessible through a default gateway, a static route must be configured on the switch in order to reach them.

See also [“Why are subnet masks important?”](#)

What should be configured on the PPTP or IPsec client?

The client should have the protocols for NetBIOS and TCP/IP configured. NetBEUI should not normally be configured.

A Windows 95 or Windows 98 client should be configured to be in the correct Workgroup for the NT domains it is trying to reach. For example, if there are domains named Engineering and Admin, and the client is to use the Engineering Domain, then it must be configured that way in its own configuration.

For PPTP only, you must also select the option Log onto Network setting under My Computer→Dial Up Networking→Connection_Name properties.

The client system’s NetBIOS name must be unique in the private network to which the client is connecting. Do not use the same name as your office desktop machine or something like “my computer.” Uniqueness is required.

What is the preferred way to access neighbors on the network?

Microsoft recommends against browsing the Network Neighborhood when tunneling. Another way to access a network resource is through the run command. For example, to access shared folders on the machine HotDog, choose Start→Run and type in \\HotDog. If you experience delays using Network Neighborhood, you may want to try this method instead.

Why should WINS settings be different for extranet access?

WINS servers cache a correspondence between IP addresses and NetBIOS names. These cached values are only invalidated by a timer and not by network activity. Therefore, a WINS server that is used heavily by clients should have its expiration timeouts set low.

In a static environment, where names and addresses correspond forever, this is not an issue. But in the extranet environment, clients are assigned new IP addresses whenever they form a tunnel. Therefore, the correspondence is transitory.

Microsoft default values for the timeouts are enormous (for example, 3 weeks). These need to be reduced for an extranet environment. Refer to [“What WINS settings are recommended?”](#) for performance tips and recommended values.

What WINS settings are recommended?

The WINS settings are available on the WINS server through the start menu, programs, Administrator Tools. Nortel Networks is currently experimenting successfully with the following values for a WINS server:

- Server Configuration
- Renewal Interval: 41 minutes
- Extinction Interval: 41 minutes
- Extinction Timeout: 24 hours
- Verify Interval: 576 hours

The Renewal Interval governs how often a client must reregister its name with the WINS server. It begins trying at one-half of the renewal interval. The Extinction Interval governs how long it is between the time a client name is released until it becoming extinct. These intervals are the most important to control when using dynamic addresses.

There is a trade-off in setting these intervals. If they are set too small there is too much additional client registration network activity. If they are set too large, transient client entries do not time-out soon enough. If you also have secondary WINS servers, the renewal interval should be the same on the secondary servers as for the primary server.

For additional information on setting interval values for a WINS configuration, refer to the Microsoft Knowledge Base article “Min. and Max. Interval Values for WINS Configuration” available at <http://support.microsoft.com/support>. A WINS server that has a heavy CPU load or network load will not perform well. To help performance:

- Do not run other intensive tasks on the WINS server.
- In the WINS configuration, disable detailed logging.
- If you have primary and secondary WINS servers, try to assign them for a balanced load.

Hosts that never change IP addresses can be given static entries in the WINS database. For example, you could configure the address of the Primary Domain Controller as static. To do this, you also have to have a statically reserved DHCP address for the Primary Domain Controller.

What can you try on the WINS server when it is not working?

You can request that the WINS server clean up its database. You can do this by going into the mappings menu item and selecting “Initiate Scavenging.”

If the database becomes very large, it can be compacted using the jetpack.exe program in \winnt\system32. Please consult the WINS Help before doing this because the server must be shut down.

In the WINS mappings entry, enter a show database command. Note the entry for `__MSBROWSE__`. This is the machine that is actually the elected master browser, and it changes frequently. If it is pointing to an invalid machine, this can cause problems.

Can I control which machine is the master browser?

When you start a computer running Windows NT Workstation or Windows NT Server, the browser service looks in the registry for the configuration parameter `MaintainServerList` to determine whether a computer will become a browser. This parameter is under:

`\HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Browser\Parameters`

For Windows 95, this parameter is under:

`HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\VNETSUP\MaintainServerList`

MaintainServerList parameter values are:

- No—This computer will never participate as a browser.
- Yes—This computer will become a browser.
- Auto—This computer, referred to as a potential browser, may or may not become a browser, depending on the number of currently active browsers.

The registry parameter IsDomainMasterBrowser impacts which servers become master browsers and backup browsers. The registry path for this parameter is \HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Browser\Parameters.

Setting the IsDomainMasterBrowser parameter entry to True or Yes makes the computer a preferred master browser. Whenever a preferred master browser starts, it forces a browser election.

When the Browser service is started on the preferred master-browser computer, the Browser service forces an election. Preferred master browsers are given priority in elections, which means that if no other condition prevents it, the preferred master browser always wins the election. This gives an administrator the ability to configure a specific computer as the master browser.

To specify a computer as the preferred master browser, set the parameter for IsDomainMasterBrowser to True or Yes in the following registry path \HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Browser\Parameters.

Unless the computer is configured as the preferred master browser, the parameter entry is always False or No. There is no user interface for making these changes; you must modify the registry.

Why are subnet masks important?

If a client does not have a WINS server or is unable to contact it, it must broadcast a query to try to locate a host. Unfortunately, Windows 95, Windows 98, and Windows NT clients do not always use the correct broadcast address when tunneling.

The following example helps explain this problem. Suppose that you are using a private net 10 address space. Assume further that you have a client with IP address 10.1.2.3 and subnet mask 255.255.0.0. This means that the net 10 space is being used like a class B address space, which is perfectly legal. The correct broadcast for this client is to 10.1.255.255. However, Microsoft clients may broadcast to 10.255.255.255, using the natural class A for net 10, in spite of their configuration.

If all hosts that the client is trying to reach lie on the same physical segment, this probably will work. This is because every host on the physical network receives the all subnets broadcast and probably respond, if appropriate.

All hosts on the segment receive the broadcast to 10.255.255.255, even if they are on different subnets (10.1.x.x. and 10.2.x.x). However, in a routed environment the situation changes. In this case, a broadcast from 10.1.2.3 to 10.255.255.255 is not forwarded to the other 10.2 subnet.

In the extranet environment you should make the remote client appear as much as possible to be on the local LAN. If the extranet host is assigned address 10.1.2.3, it should behave as if it is on the 10.1 LAN.

When 10.1.2.3 broadcasts to find a network neighbor, it (incorrectly) sends to 10.255.255.255. Normal routing functionality would not forward such a packet. The switch finds the best match among its physical interfaces (10.1 in this case) and modify the broadcast to be correct for that interface (10.1.255.255 here).

In this example, if the switch's 10.1 interface had been configured with any subnet mask other than 255.255.0.0, the broadcast would not have been converted as desired.

What should I do about subnets?

Every private interface on the switch should be configured to have the same subnet mask as all of the clients residing on that subnet.

Why is there a delay in discovering the Network Neighborhood (with tunnels)?

NetBIOS treats the modem interface as if it is two different interfaces: the original modem and the tunnel. It designates the original modem as the primary interface. (You can observe this by typing `route print` in a DOS command shell.) If you tunnel over a LAN instead of a modem, the LAN adapter is designated as the primary interface.

When first instructed to seek the Network Neighborhood, NetBIOS always tries the primary interface first. This is always the wrong choice because NetBIOS always tries to send using the IP address assigned by the ISP (or possibly the address of another adapter) instead of the address assigned to the tunnel by the switch.

The outcome is somewhat different for IPsec and PPTP. For IPsec, the switch client recognizes this incorrect behavior and refuses to even send the packets. You can see a counter of the number of invalid packets of this type on the client under status, Invalid IP address.

With PPTP, the client does send the packets, but they are rejected at the switch as invalid tunneled packets because the source address does not match the switch-assigned address. If you inspect the event log, there are messages of the form, Bad source address in tunnel, and the session/details counter for source address drops increases.

After about 10 to 15 seconds, NetBIOS gives up on the primary interface and switches to the correct tunnel interface and proceeds browsing the Network Neighborhood.

Why can't I browse another client in a different tunnel?

Cause: If you are not using a WINS server, this is not possible because network browsing requires broadcasts from one tunnel to another.

Action: Use a WINS server to browse another client in a different tunnel. When the clients tunnel in, they should register with the WINS server. Be sure that the client you want to browse has enabled the Log onto Network setting under My Computer→Dial Up Networking→Connection_Name properties.

Where can I get more information on troubleshooting dial-up connections?

The Microsoft Knowledge Base article “Dial-Up Networking 1.2 Dun12.doc” file, available from <http://support.microsoft.com/support>, contains help for resolving common dial-up problems:

Depending on the service provider, some POPs may not support LCP Options. If your connection constantly gets declined after the modems synchronize, and you know your password is correct, try disabling this option. The Microsoft Knowledge Base “Service Pack 2 May Cause Loss of Connectivity in Remote Access” article contains more details.

Where can I get more information on configuring PPTP on my client?

There are many articles in the Microsoft Knowledge Base on configuring PPTP for Windows NT, Windows 98, and Windows 95. Refer to the section “[Additional information](#)” for a partial list. In addition, Microsoft has the following white papers available at <http://support.microsoft.com/support> that contain helpful information:

- Microsoft Windows 95/Windows NT White Paper, “Installing, Configuring, and Using PPTP with Microsoft Clients and Servers”
- Microsoft Windows NT Server White Paper, “Understanding PPTP”

You must create a connection definition for your initial Internet link via your service provider. A separate connection definition is needed for creating the PPTP tunnel. A common configuration problem experienced during initial PPTP setup is the failure to select the PPTP VPN adapter (instead of the modem) on the PPTP connection definition in Dialup Networking.

What DNS and WINS servers should I set for the dial-up connection?

There should be no need to set these servers statically on your dial-up client as this information is dynamically downloaded from the switch for PPTP, IPsec, and L2F tunnels at connect time.

Why does DNS resolve hosts to different addresses when a tunnel connection is active?

Cause: When a tunnel connection is activated, additional DNS servers are downloaded from the extranet device (for example, the switch) to your client. In the case of Microsoft Windows 95, Windows 98, and Windows NT operating systems, the new DNS servers are added to the list of DNS servers that were assigned by your ISP. This applies to PPTP as well as IPsec tunnels. In general, the DNS servers downloaded by the extranet device provide host-name-to-address translation for hosts within a private network while the ISP-based DNS servers can be expected to translate public host names.

For Windows 95/98 and Windows NT, when a host name must be translated to an IP address (for example to browse the Web or get e-mail), all DNS servers are queried in a shotgun style. The first server to respond with an IP address wins. This can produce some interesting behavior if a host name resolves to one address on the private network and another on the public Internet. For example, you have host mail.mycompany.com. Internally this host resolves to 10.0.0.282 and externally to 146.113.64.231.

Action: To avoid problems when using a mixture of internal and external DNS services, it is essential to avoid using names that can resolve to different addresses. In the preceding example, host 10.0.0.282 should probably be renamed pop.mycompany.com. Then users can be informed to use the hostname pop.mycompany.com to retrieve electronic mail whether in the office or connected via a tunnel link. The original retail release of Windows 95 requires the Winsock DNS Update (wsockupd) to properly function with multiple DNS servers.

My downloaded DNS servers for my tunnel connection do not work

Cause: The Microsoft Windows 95/98 and Windows NT operating systems attempt to ping new DNS servers before adding them to the current list of servers.

Action: As a quick test, try to ping (with the tunnel connection active) the DNS servers that the extranet device is downloading at tunnel startup. If you cannot ping the servers, a basic connectivity problem using the tunnel connection exists.

To view the current list of DNS servers at any time use the MS-DOS command `ipconfig/all` on Windows NT or `winipcfg` on Windows 95 or Windows 98.

Why, after disconnecting a PPTP tunnel, do I get an immediate error reconnecting?

Cause: After disconnecting a PPTP tunnel then immediately trying to reconnect, the PPTP client indicates the connection is busy or otherwise unavailable. On Windows 95 this is caused by the PPTP control channel socket being improperly shut-down by the client.

Action: You can wait for the socket to time out, but it is often more expedient to reboot. On Windows NT a similar problem can be encountered, but caused by a TCP checksum error generated by the Microsoft IP stack. The only resolution for the Windows NT error condition as of this writing is to reboot.

Additional information

Below is a list of some of the Microsoft Knowledge Base topics you can browse for information related to dial-up and tunnel configuration. To view these topics, go to <http://support.microsoft.com/support>. Use the Search Support Online feature to search on the title you want.

- Troubleshooting Internet Service Provider Login Problems
- Service Pack 2 May Cause Loss of Connectivity in Remote Access
- Troubleshooting Modem Problems Under Windows NT 4.0
- Dial-Up Networking 1.2 Dun12.doc File (Windows 95 PPTP Troubleshooting)
- How to Troubleshoot TCP/IP Connectivity with Windows NT
- Remote Access Service (RAS) Error Code List for Windows NT 4.0
- RAS Error 720 When Dialing Out
- Troubleshooting PPTP Connectivity Issues in Windows NT 4.0
- PPTP Registry Entries
- Connecting to Network Resources from Multihomed Computer
- How to Force 128-bit Data Encryption for RAS
- Login Validation Fails Using Domain Name Server

Solving general problems

This section contains general recommendations and explains some common problems that can occur with common Web browsers, the Nortel Networks Contivity VPN Client Manager, and the switch.

Web browser problems and the Contivity VPN Client Manager

If you have a problem browsing the Nortel Networks Contivity VPN Client Manager, start by checking the following recommendations to ensure that you are using the correct Web browser version and settings. For additional troubleshooting, check the described Web browser problems and solutions, error messages, and tips described later in this section.

Nortel Networks Contivity VPN Client Manager uses Java* and HTML features. For the management interface to function properly, verify that your Web browser meets the following minimum requirements:

- Platforms supported—Windows 95, Windows 98, Windows NT, or Macintosh*.
- Display setting—256 colors or greater.
- Browser version—Microsoft Internet Explorer*, Version 4.0 or later. Netscape Communicator*, Version 4.0 or later. Not using a recent version of Internet Explorer causes the upper-left and top-left corners of the management screens to remain gray rather than displaying the navigational menu and the current menu selection, respectively.
- ActiveX Scripts*, Java*, and JavaScript*—You must Enable both ActiveX and Java programs in Internet Explorer, and enable both Java and JavaScript in Netscape* Communicator* for proper switch Web management screens. These options are enabled by default on both Web browsers.

Enabling Web browser options

To make sure these options are enabled in Internet Explorer, from the Internet Explorer menu bar, choose View→Options→Security and choose:

- Run ActiveX scripts—If this option is disabled, navigational titles are not updated, and the Logoff and Help buttons do not work.
- Enable Java programs—If this option is disabled, navigational menus do not appear.

To make sure these options are enabled in Netscape, from the Netscape menu bar, choose Edit→Preferences→Advanced, and choose:

- Enable Java—If this option is disabled, navigational menus do not appear.
- Enable JavaScript—If this option is disabled, navigational titles are not updated, and the Logoff and Help buttons do not work.

Long delays when Web browsing

Cause: HTTP. Sometimes while using the HTTP Web interface you can experience long delays (greater than five minutes).

Action: Wait until the requested screen is fully delivered before clicking a new screen request.

Improving performance with Internet Explorer 4.0

Nortel Networks recommends that you create a DNS server entry for your Management IP address. This alleviates a noticeable delay in loading the initial Main Menu and navigational screens.

Clearing your Web browser cache when upgrading

To avoid problems when upgrading software revision levels (for example, moving from Version 01_01.16 or V01_05_01.28 to V01_00.33), Nortel Networks recommends that you clear your browser cache and exit the browser and all associated windows (such as mail and news readers). Refer to the following section for browser cache clearing instructions.

Clearing cache

Browsers cache pages that they read to improve performance if the same page is requested again. The switch HTTP server allows the browsers to cache the Java class files, and all image files, but not to cache the body pages that contain the dynamically generated information. Both Internet Explorer and Netscape allow you to clear the browser cache which causes all pages to be re-requested the next time they are required. To manually clear the browser cache in Internet Explorer V4.x, select View→Internet Options, and click Delete Files. To manually clear the browser cache in Netscape V4.x, go to Edit→Preferences→Advanced→Cache and select Clear disk and memory cache.

Web browser error messages

No data in post message

Cause: This message often appears in the main body page if you use the browser's back arrow to revisit a previously displayed page. The browser displays this message when it knows you are revisiting a dynamically generated page.

Action: To see the page, you must use the left navigational area to select it.

Internal error message

Cause: The HTTP server was unable to allocate memory. This indicates that the switch is very low on memory.

Action: Terminate any unnecessary tasks to free up memory. It may be necessary to reboot the switch. If this condition reoccurs, there may be a serious problem. Contact Nortel Networks Customer Support.

Document not found message

Cause: This message is returned when the HTTP server cannot find the requested page on the switch. This could happen because the Java navigation index file is out of synch with the rest of the system. A corrupted or incorrectly cached index file could also cause this problem.

Action: Clear your browser cache or restart your browser to correct this problem.

New administrator login ignored

Cause: Internet Explorer saves in its cache your User ID and Password and automatically resends those values on subsequent login attempts. Therefore, when prompted after an Idle Timeout, the User ID and Password value you enter are ignored, and Internet Explorer sends the original User ID and Password. For example, if you log in as Administrator with password abc123De, and you subsequently log out. If you log in again, this time as DottieDoe with password FGh45678, in spite of the different login and password, Internet Explorer sends Administrator with passwordabc123De.

Action: When you log off the switch, close out of the Web browser completely (shut down the browser). This clears the cache and the next time that you log in you are starting fresh.

Excess resource consumption using Internet Explorer

Cause: Internet Explorer has a known problem with excessive memory consumption using Java applets. Over time, this problem can cause serious overall system performance degradation.

Action: If you notice that your system's performance seems to slow down for no reason, close and restart Internet Explorer. This releases unused memory and should improve system performance. Go to <http://premium.microsoft.com/support/kb/articles/q173/1/45.asp> for details.

Internet Explorer 4.0 multiple Help windows

Cause: In Internet Explorer 4.0, if you select context-sensitive Help and do not close the Help window after viewing, you might end up with multiple Help windows open.

Action: Close Help windows after viewing them.

Distorted background images

Cause: In Netscape versions prior to 4.0 where you configure your Windows 95, Windows 98, or Windows NT system for 8-bit color (256 colors or less), images might appear distorted in the navigational area.

Action: To avoid this situation, increase the color display setting to 256 or greater. You should check with your video card manufacturer's documentation to confirm that your video card supports 256 colors or greater.

Reporting a problem with a Web browser

When reporting a problem with a browser to Nortel Networks, include the following information:

- Workstation operating system and version
- Browser vendor and version (major and minor version)
- Cache setting (size in Netscape, percent of drive for Internet Explorer)
- Verify document setting (every time or once per session)

Switch problems

Excessive active sessions logged

Cause: The number of Active Sessions can reach in excess of 4 billion. This is an erroneous number that results from a negative number of sessions that can appear.

Action: Restart the system.

Power failure

Cause: The power supplies can become unseated during shipping. When this problem occurs, you might not be able to start the switch, or a Warning might be posted to the Status→Health Check screen indicating a potential problem.

Action: If necessary, remove the front bezel as directed in the Getting Started Guide, then push the bottom of the power supply in to reseal it.

Cannot convert from an internal address pool to an external DHCP server

Cause: You cannot convert IP address distribution from an Internal Address Pool to an External DHCP server while sessions are active.

Action: Go to the Admin→Shutdown screen, and select Disable Logins after Restart. After everyone has logged off of the switch, then you can convert from an Internal Address Pool to an External DHCP server.

Group and user profile settings not saved

Cause: Saving a configuration from the Admin→Configs screen Save Current Configurations option saves only the operational parameters in the configuration file, such as interface IP addresses and subnet masks, backup host IP addresses, DNS names.

Action: To completely back up the switch's configuration, you must also back up the LDAP database, which contains the group and user profiles, filters, backup file names.

- 1 Go to the Servers→LDAP screen and click Stop Server.
- 2 Enter a file name in the Backup/Restore LDAP Database field. The name should conform to the MS-DOS naming convention and append the filename with LDF (for example, ldapone.ldf). The restore process can take anywhere from five minutes for a very small LDAP database to several hours for a very large database.
- 3 You can view the progress of the restoration from the Admin→Health Check screen.

Upgrade fails with large log files

Cause: Currently there is a problem with large Log files that can cause the Upgrade process to fail.

Action: Before performing an upgrade, Nortel Networks recommends that you delete the Account Log and Log files:

- 1 Choose Admin→File System, click on ACCTLOG, and double-click Details. Under Action, double-click Delete Directory. A confirmation dialog box requests that you verify the deletion. Click OK. Do not select Remove Directory Name.
- 2 Choose Admin→File System, click on LOG, and double-click Details. Under Action, double-click Delete Directory. A confirmation dialog box requests

that you verify the deletion. Click OK. Do not select Remove Directory Name.

Restart fails after using recovery and reformatting the hard disk

Cause: When using the recovery disk and reformatting the hard disk, sometimes the system will not restart.

Action: Power cycle the system using the green power button on the back of the switch.

Solving routing problems

The following sections describe routing problems.

Client address redistribution problems

The number of current Utunnel host users may display more than the configured maximum.

Cause: This is not an error and is the running state of the system. For example, if you configured a maximum of 200 and have 150 logins, the screen will display the maximum as 200 and the current as 150. If you then modify the maximum to 100, the screen will display the maximum as 100 and the current as 150. As users log out, the current number will eventually be no greater than the maximum.

Action: No action.

Client address redistribution is enabled and the client is logged in, but the client is not communicating with the private network.

Cause: Client address redistribution is not enabled.

Action: Have the client log in again. Client address redistribution only takes effect if the client logs in when it is enabled.

- 1** Check the Routing→Policy screen to be sure Utunnel routes are enabled.
- 2** Check to be sure that the routing protocols OSPF and RIP are properly set up.
- 3** Check to be sure you have the correct address ranges if you configured summarization.
- 4** Be sure you have an Advanced Routing license if you are using OSPF for client address redistribution.

Appendix A

MIB Support

The Contivity VPN Switch supports the management information base (MIB) for use with network management protocols in TCP/IP-based Internets and TCP/IPX-based networks. The switch only supports SNMP Gets. It does not support SNMP Sets.

Nortel Networks also provides proprietary MIBs for the switch's SNMP trap support. The MIBs, `cestraps.mib` and `newoak.mib`, are available on the Contivity VPN Switch distribution CD in the `Doc` directory.

SNMP RFC support

This section discusses the SNMP-related RFCs that the Contivity VPN Switch supports.

Novell IPX MIB

The switch supports the IPX MIB that is distributed by Novell, Inc.

Novell RIP-SAP MIB

The switch supports the IPX RIP-SAP MIB that is distributed by Novell, Inc.

RFC 1850 -- OSPF Version 2 Management Information Base

The switch supports RFC 1850, *OSPF Version 2 Management Information Base*. As stated in the introduction to the RFC, the RFC “defines a portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing the Open Shortest Path First Routing Protocol.”

RFC 1724 -- RIP Version 2 MIB Extension

The switch supports RFC 1724, *RIP Version 2 MIB Extension*. As stated in the introduction to the RFC, the RFC “defines a portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines the objects for managing RIP Version 2.”

RFC 1213 -- Network Management of TCP/IP-Based Internets MIB

The switch supports RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB II*. This RFC provides the architecture and system for managing TCP/IP-based internets. With the exception of the EGP Group (Section 6.10) and the Transmission Group (Section 6.11), the switch provides full support for the RFC.

RFC 2667 -- IP Tunnel MIB

The switch supports RFC 2667, *IP Tunnel MIB*. As stated in the introduction to the RFC, it “describes a Management Information Base (MIB) used for managing tunnels of any type over IPv4 networks, including GRE [16,17], IP-in-IP [18], Minimal Encapsulation [19], L2TP [20], PPTP [21], L2F [25], UDP (e.g., [26]), ATMP [22], and IPv6-in-IPv4 [27] tunnels.”

RFC 2787 -- VRRP MIB

The switch supports RFC 2787, *Definitions of Managed Objects for the Virtual*

Router Redundancy Protocol. As stated in the introduction to the RFC, it “defines an extension to the Management Information Base (MIB) for use with SNMP-based network management. In particular, it defines objects for configuring, monitoring, and controlling routers that employ the Virtual Router Redundancy Protocol (VRRP).”

RFC 2737 -- Entity MIB

This MIB contains five tables, we have partially implemented two of these tables.

```
*entPhysicalTable
entLogicalTable
entLPMappingTable
*entAliasMappingTable
entPhysicalContainsTable
```

The entPhysicalTable provides a listing of the hardware elements that are present in the system. For example each slot is listed, and if there is a card in the slot, then the card and any ports on the card. The exception to this is the hardware accelerator which does not appear in the table). The listing shows element relationships via the columns entPhysicalContainedIn and entPhysicalParentRelPos. The only columns that have been implemented are:

```
entPhysicalIndex
entPhysicalDescr (although the value is not strictly what the MIB
specifies)
entPhysicalContainedIn
entPhysicalClass
entPhysicalParentRelPos
entPhysicalName
entPhysicalIsFRU
```

All other columns will return an appropriate default value for the object.

The entAliasMappingTable provides a mapping from entPhysicalIndex to ifTable.ifIndex. Hence by walking this table, a management station can determine the ifIndex associated with a physical port.

RFC 1573 -- IanaIfType MIB

This MIB contains the enumerations for rfc2233 ifTable.ifType. These enumerations describe the various types of interfaces that ifTable can support.

RFC 2233 -- If MIB

This MIB is the latest evolution of rfc1213 Interfaces group, plus several new objects.

RFC 2571-- Snmp-Framework MIB

This MIB provides textual conventions and object definitions used in the SNMP agent architecture.

CES MIB

This MIB contains CES proprietary MIB data. For instance the “ping MIB” is contained in this file. The ping MIB, via an SNMP GET REQUEST, causes the CES to ping another device and get statistics based on the results of the ping. For instance sending a PDU specifying pingAverageTime.192.32.250.248.4.4076, sends four pings, of 4076 bytes, to address 192.32.250.248. (It actually sends five because one ping is sent by itself so that if the device being pinged is the other end of a Branch Office tunnel, it ensures that the tunnel is brought up before trying to send pings through the tunnel. This ping is not counted in the statistics.) The object returns the values of:

```
-2 Invalid parameter(indices).  
-1 No reply.  
0 Less than 16ms average time.  
>0 The average time.
```

The objects and their parameters(indices) are:

```
pingAverageTime - returns the average ping time for the set of  
specified pings.  
pingPercentLoss - returns the percentage of loss.
```

The first index is the IP address to ping. The second index is the number of pings, if this is not specified or is an invalid value it defaults to 3. The third index is the size of the ping request. If it is not specified or is an invalid value then it defaults to 1024.

cestraps.mib -- Nortel Networks proprietary MIB

This section lists the contents of the cestraps.mib, the Nortel Networks MIB for the Contivity VPN Switch.

```
-- Each Trap contains the Trap OID as well as the following OIDs:
--     SeverityLevel
--     System Name
--     System Date
--     System Time
--     System Uptime
--
NEWOAKTRAP DEFINITIONS ::= BEGIN

IMPORTS
    enterprises                FROM RFC1155-SMI
    DisplayString              FROM RFC1213-MIB
    OBJECT-TYPE                FROM RFC-1212
    TRAP-TYPE                  FROM RFC-1215;

-- This MIB module uses the extended OBJECT-TYPE macro as
-- defined in [9], and the TRAP-TYPE macro as defined in [10].

contivity                     OBJECT IDENTIFIER ::= { enterprises 2505 }

ContivitySnmpTraps OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "Nortel Networks Inc's Enterprise trap."
    ::= { contivity 1 }
```

newoak.mib

This section provides the contents of the newoak.mib, which defines the “newoak” enterprise ID, the “contivity” object identifier, and the sysObjectIDs for each of the switch models.

```
-- This MIB module uses the extended OBJECT-TYPE macro as
-- defined in [9], and the TRAP-TYPE macro as defined in
[10].

newoak    OBJECT IDENTIFIER ::= { enterprises 2505 }

-- The following MODULE-IDENTITY definition can be commented out if
the MIB parser
-- you are using has trouble parsing it. If you do comment it out,
then uncomment
-- the following object identifier definition.
--   contivity OBJECT IDENTIFIER ::= {newoak 1}
--
contivity MODULE-IDENTITY
    LAST-UPDATED "0004252130Z" -- April 25, 2000 7:30pm EST
    ORGANIZATION "Nortel Networks, Inc."
    CONTACT-INFO
        "support@nortelnetworks.com
        Postal: Nortel Networks, Inc.
              80 Central St.
              Boxboro, MA 01719
        Tel:    +1 978 264 7100
        E-Mail: support@nortelnetworks.com"

    DESCRIPTION
        "This MIB defines the sysObjectIDs for different
        variations of the Convitiy Extranet Switch."
        ::= { newoak 1 }
-- IDENTIFIER ::= {newoak 1}
contivityExtranetSwitch2000 OBJECT IDENTIFIER ::= {newoak 2}
contivityExtranetSwitch1000 OBJECT IDENTIFIER ::= {newoak 3}
contivityExtranetSwitch4500 OBJECT IDENTIFIER ::= {newoak 4}
contivityExtranetSwitch15XX OBJECT IDENTIFIER ::= {newoak 5}
contivityExtranetSwitch2500 OBJECT IDENTIFIER ::= {newoak 6}
contivityExtranetSwitch2600 OBJECT IDENTIFIER ::= {newoak 7}
contivityExtranetSwitch1600 OBJECT IDENTIFIER ::= {newoak 8}
contivityExtranetSwitch4600 OBJECT IDENTIFIER ::= {newoak 9}

END
```


Hardware-related traps

```

hardwareTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 1}

-- Trap #1001
hardDisk1Status OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Hard Disk Number 1 Status."
    ::= {hardwareTrapInfo 1}

-- Trap #1002
hardDisk0Status OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Hard Disk Number 0 Status."
    ::= {hardwareTrapInfo 2}

-- Trap #1003
memoryUsage OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Memory Usage Status."
    ::= {hardwareTrapInfo 3}

-- Trap #1004
LANcardStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of any LAN cards on the system."
    ::= {hardwareTrapInfo 4}

-- Trap #1005
CPUtwoStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of second CPU."
    ::= {hardwareTrapInfo 5}

-- Trap #1006
fanOneStatus OBJECT-TYPE
    SYNTAX  DisplayString

```

```
        ACCESS    read-only
        STATUS    mandatory
        DESCRIPTION "Status of the first CPU fan."
        ::= {hardwareTrapInfo 6}

-- Trap #1007
fanTwoStatus OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION "Status of the second CPU fan."
    ::= {hardwareTrapInfo 7}

-- Trap #1008
chassisFanStatus OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION "Status of the Chassis fan."
    ::= {hardwareTrapInfo 8}

-- Trap #1009
fiveVoltsPositive OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION "Status of +5 Volt power."
    ::= {hardwareTrapInfo 9}

-- Trap #10010
fiveVoltsMinus OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION "Status of -5 Volt power."
    ::= {hardwareTrapInfo 10}

-- Trap #10011
threeVoltsPositive OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION "Status of +3 Volt power."
    ::= {hardwareTrapInfo 11}

-- Trap #10012
twoDotFiveVA OBJECT-TYPE
    SYNTAX      DisplayString
```

```
ACCESS read-only
STATUS mandatory
DESCRIPTION "Status of 2.5VA power."
::= {hardwareTrapInfo 12}

-- Trap #10013
twoDotFiveVB OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of 2.5VB power."
    ::= {hardwareTrapInfo 13}

-- Trap #10014
twelveVoltsPositive OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of +12 Volt power."
    ::= {hardwareTrapInfo 14}

-- Trap #10015
twelveVoltsMinus OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of -12 Volt power."
    ::= {hardwareTrapInfo 15}

-- Trap #10016
normalTemperature OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of normal temperature reading."
    ::= {hardwareTrapInfo 16}

-- Trap #10017
criticalTemperature OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of critical temperature reading."
    ::= {hardwareTrapInfo 17}

-- Trap #10018
chassisIntrusion OBJECT-TYPE
    SYNTAX DisplayString
```

```
        ACCESS read-only
        STATUS mandatory
        DESCRIPTION "The chassis intrusion sensor indicates that
                     the unit has been opened."
        ::= {hardwareTrapInfo 18}

-- Trap #10019
dualPowerSupply OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of the redundant power supplies."
    ::= {hardwareTrapInfo 19}

-- Trap #10020
t1WANStatus OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of T1 WAN card(s)."
    ::= {hardwareTrapInfo 20}

-- Trap #10021
t3WANStatus OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of T3 WAN card(s)."
    ::= {hardwareTrapInfo 21}
```

Server-related traps

```

serverTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 2}

-- Trap #3001
radiusAcctServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of External Radius Accounting Server."
    ::= {serverTrapInfo 1}

-- Trap #3002
backupServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of External Disk Backup Server."
    ::= {serverTrapInfo 2}

-- Trap #3003
diskRedundancy OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of Local Disk Redundancy."
    ::= {serverTrapInfo 3}

-- Trap #3004
IntLDAPServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of Internal LDAP Server."
    ::= {serverTrapInfo 4}

-- Trap #3005
LoadBalancingServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of Load Balancing Server."
    ::= {serverTrapInfo 5}

-- Trap #3006
DNSServer OBJECT-TYPE
    SYNTAX  DisplayString

```

```
        ACCESS read-only
        STATUS mandatory
        DESCRIPTION "Status of DNS Server."
        ::= {serverTrapInfo 6}

-- Trap #3007
SNMPServer OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of SNMP Server."
    ::= {serverTrapInfo 7}

-- Trap #3008
IPAddressPool OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of the IP address pool."
    ::= {serverTrapInfo 8}

-- Trap #3009
ExtLDAPServer OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of External LDAP Server."
    ::= {serverTrapInfo 9}

-- Trap #30010
radiusAuthServer OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of Radius Authentication Server."
    ::= {serverTrapInfo 10}

-- Trap #30011
certificateServer OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Status of Certificates Validity."
    ::= {serverCESTrapInfo 11}
```

Software-related traps

```
softwareTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 3}

-- Trap #5001
NetBuffers OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Network buffer usage."
    ::= {softwareTrapInfo 1}

-- Trap #5002
fireWall OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of internal firewall."
    ::= {softwareTrapInfo 2}
```

Login-related traps

```
loginTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 4}

-- Trap #101
failedLogin OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Failed Login Attempt."
    ::= {loginTrapInfo 1}
```

Intrusion-related traps

```
intrusionTrapInfo OBJECT IDENTIFIER
 ::= {ContivitySnmpTraps 5}

-- Trap #201
securityIntrusion OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Login Security Intrusion."
    ::= {intrusionTrapInfo 1}
```

System-related traps

```
-- Trap #401
powerUpTrap OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Power Up."
    ::= {ContivitySnmpTraps 6}

-- Trap #601
periodicHeartbeat OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Periodic Heartbeat."
    ::= {ContivitySnmpTraps 12}
```


Information passed with every trap

```

SeverityLevel OBJECT-TYPE
    SYNTAX INTEGER
    {
        fatal(1),
        major(2),
        minor(3),
        informational(4),
        insignificant(5),
        reversal(6)
    }
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "Severity of specific trap."
    ::= {ContivitySnmpTraps 7}

systemName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "System Name."
    ::= {ContivitySnmpTraps 8}

systemDate OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "System Date."
    ::= {ContivitySnmpTraps 9}

systemTime OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "System Time."
    ::= {ContivitySnmpTraps 10}

systemUpTime OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "System Up Time."
    ::= {ContivitySnmpTraps 11}

```

The following table provides explanations for the traps.

Table 19 Trap categories

Hardware	
1.3.6.1.4.1.2505.1.1.0.1001	hardDisk1StatusTrap
1.3.6.1.4.1.2505.1.1.0.1002	hardDisk0StatusTrap
1.3.6.1.4.1.2505.1.1.0.1003	memoryUsageTrap
1.3.6.1.4.1.2505.1.1.0.1004	lanCardStatusTrap
1.3.6.1.4.1.2505.1.1.0.1005	cpuTwoStatusTrap
1.3.6.1.4.1.2505.1.1.0.1006	fanOneStatusTrap
1.3.6.1.4.1.2505.1.1.0.1007	fanTwoStatusTrap
1.3.6.1.4.1.2505.1.1.0.1008	chassisFanStatusTrap
1.3.6.1.4.1.2505.1.1.0.1009	fiveVoltsPosStatusTrap
1.3.6.1.4.1.2505.1.1.0.10010	fiveVoltsMinusTrap
1.3.6.1.4.1.2505.1.1.0.10011	threeVoltsPositiveTrap
1.3.6.1.4.1.2505.1.1.0.10012	twoDotFiveVATrap
1.3.6.1.4.1.2505.1.1.0.10013	twoDotFiveVBTrap
1.3.6.1.4.1.2505.1.1.0.10014	twelveVoltsPositveTrap
1.3.6.1.4.1.2505.1.1.0.10015	twelveVoltsMinsTrap
1.3.6.1.4.1.2505.1.1.0.10016	normalTemperatureTrap
1.3.6.1.4.1.2505.1.1.0.10017	criticalTemperatureTrap
1.3.6.1.4.1.2505.1.1.0.10018	chassisIntrusionTrap
1.3.6.1.4.1.2505.1.1.0.10019	dualPowerSupplyTrap
1.3.6.1.4.1.2505.1.1.0.10020	t1WANStatusTrap
1.3.6.1.4.1.2505.1.1.0.10021	t3WANStatusTrap
1.3.6.1.4.1.2505.1.1.0.10022	hwAccelTrap
Server	
1.3.6.1.4.1.2505.1.2.0.3001	radiusAcctServerTrap
1.3.6.1.4.1.2505.1.2.0.3002	backupServerTrap
1.3.6.1.4.1.2505.1.2.0.3003	diskRedundencyTrap
1.3.6.1.4.1.2505.1.2.0.3004	intLDAPServerTrap
1.3.6.1.4.1.2505.1.2.0.3005	loadBalancingServerTrap
1.3.6.1.4.1.2505.1.2.0.3006	dnsServerTrap

Table 19 Trap categories

Server	
1.3.6.1.4.1.2505.1.2.0.3007	snmpServerTrap
1.3.6.1.4.1.2505.1.2.0.3008	ipAddressPoolTrap
1.3.6.1.4.1.2505.1.2.0.3009	extLDAPServerTrap
1.3.6.1.4.1.2505.1.2.0.30010	radiusAuthServerTrap
1.3.6.1.4.1.2505.1.2.0.30011	certificateServerTrap
Software	
1.3.6.1.4.1.2505.1.3.0.5001	netBuffersTrap
1.3.6.1.4.1.2505.1.3.0.5002	FireWallTrap
1.3.6.1.4.1.2505.1.3.0.5003	FipsStatusTrap
Failed Login	
1.3.6.1.4.1.2505.1.4.0.101	FailedLoginTrap
Intrusion	
1.3.6.1.4.1.2505.1.5.0.201	SecurityIntrusionTrap
Presence	
1.3.6.1.4.1.2505.1.0.401	PowerUpTrapEntry
1.3.6.1.4.1.2505.1.0.601	PeriodicHeartbeatTrap

Appendix B

Using Serial PPP

The Serial Point-to-Point Protocol (PPP) feature allows you to manage the switch from a remote location using PPP and the serial interface. If the switch were to become unreachable over the Internet, you could still dial up and manage it through the serial interface menu.

With this feature, the serial interface becomes much like a private WAN interface. You can manage through it or even tunnel through it. You can enable Serial PPP support on the switch using the Web interface (System→Settings). When configuring Serial PPP, you can set the switch to Auto Detect, or you can specify that either PPP or the Serial Menu are the options available through the serial port.

Serial PPP authentication is performed by the Password Authentication Protocol (PAP), which uses a standard user ID and a password that is sent in the clear. When authenticated, the serial interface acts like a private WAN interface.

Establishing a serial PPP connection

To enable Serial PPP:

- 1 Set up a Dial-Up Networking connection.
- 2 Set up the modem.
- 3 Set up the switch.
- 4 Dial into the switch using the Primary Administrator's user name and password.

Setting up a Dial-Up Networking connection

To establish a Serial PPP connection using a Microsoft Dial-Up Networking connection from the client system:

- 1** Double-click on My Computer.
- 2** Double-click on the Microsoft Dial-Up Networking icon.
- 3** Set the COM port baud rate on the client system to be compatible with the switch's baud rate. The rates should be the same to establish a connection. Possible rates are:
 - 9600 (default)
 - 19200
 - 38400
 - 56000
- 4** Go to Server Types, and under Type of Dial-Up Server, select PPP: Internet, Windows NT Server, Windows 95. Make sure that none of the Advanced options are set.
- 5** Go to Allowed network protocols, and select TCP/IP.
- 6** Go to TCP/IP Settings, and specify your IP address. This is the Management IP address that the switch uses to communicate with the client that is dialing in through the modem.
- 7** Click Server Assigned name server addresses.
- 8** Unclick IP header compression.
- 9** Click Use default gateway on remote network.
- 10** Do not configure Scripting and Multilink.
- 11** Click Configure the client modem, and use the following settings:
 - 8 data bits
 - 1 stop bit
 - No parity
 - Hardware flow control

Do not choose Log On to Network if the selection appears.

Setting up the modem

The following procedure assumes you are using a 3Com/US Robotics 56K x2 modem. It describes how to set up a modem to communicate with the switch using a Dial-Up Networking connection.

Table 20 DIP switches configuration

Parameter	Setting
Data Terminal Ready	On
Verbal Result Codes	On
Suppress Result Codes	On
Echo Offline Commands	Off
Auto Answer (must be set)	On
Carrier Detect Normal	On
Load NVRAM Defaults	On
Dumb Mode	Off

Setting up the switch

To set up the switch's parameters through the Web interface (System→Settings):

- 1 Go to the System→System Settings screen and select one of the following modes of operation under the Serial Port option:
 - **Serial Menu (default)**--leaves the switch's serial interface in the traditional serial menu mode. In this mode, no serial PPP is supported. When connecting a program such as Hyper Terminal to the interface, the standard serial interface menu appears. In Auto Detect Mode, if you are using a terminal emulator, such as Hyper Terminal, you must press Enter several times to get the logon and password prompt. Also, you can ignore the modem initialization string (which might not be in use) that is displayed on the Hyper Terminal screen.
 - **PPP**--you can set up the switch to use the Point-to-Point Protocol (PPP) protocol over the serial port. This feature allows you to manage the switch from a remote location using PPP and the serial interface. If the switch were to become unreachable over the internet, you could still dial up and manage it through the serial interface menu. This feature allows you to

access all of the management services (HTTP, Telnet, FTP, SNMP) through the Web interface. Once a session is established through PPP, the serial interface acts as a private WAN interface with an internal IP address (0.0.1.35).

- Auto detect--automatically detects whether the connected device is using PPP or serial menu mode at startup. The switch cannot determine the device's baud rate, nor can it determine a change from PPP to serial menu mode, except upon startup. Auto Detect checks the mode each time the switch is restarted. When performing its Auto Detect check, the switch sends out AT command set characters to configure a modem if one is attached.

When the switch is in Auto Detect mode, and if a terminal session is connected and the terminal baud rate is the same as the switch's, the terminal displays the AT command sets on the screen. Simply press Enter several times until a serial menu session starts. You should use the Auto Detect Mode rather than PPP Mode. Using PPP mode could leave the switch in a state such that you could never manage it from the serial interface menu directly. If this happened, you would still have to manage the switch through a PPP application (such as Dial-Up Networking). Directly connecting a serial cable and running Hyper Terminal would not work because the interface would only recognize PPP.

- 2 Select one of the following Baud Rates to match the baud rate of your terminal. After you select the baud rate, you must click the Reset button to change the port to the selected baud rate. This option is necessary for PPP if a modem initialization string specifies a fixed baud rate.

- 57600
- 38400
- 19200
- 9600 (default)

- 3 Enter the modem initialization string. Refer to the manufacturer's documentation to learn the vendor-specific character initialization string. Preconfiguring the modem and using the switch's default initialization string (ATZ) provides the best results.

A sample 3Com/US Robotics 56K modem initialization string that instructs the external modem to connect at 19,200 Kb/s is ATZAT&B1AT&N10.

- 4 Click the Reset button to reset the port to the selected baud rate and apply any other modem changes.

Dialing in to the switch

Use the standard dial-up networking procedure to connect to the switch. After connecting, you can then manage the switch using either Telnet (for the command line interface) or the browser-based GUI. Use the switch's management IP address for the telnet session or the browser's destination URL.

Troubleshooting Serial PPP

When the serial port is set up for PPP only, you can still do in-band Web management.

Cause:

I have a modem connected, but I cannot get a PPP connection.

Actions:

- Verify that the modem supports the switch's selected baud rate. Most connection problems occur because the modem is not operating at the same baud rate as the switch. For example, a 3Com/US Robotics 56 Kb/s modem's default baud rate when attempting to establish a connection to the switch is 38400, but the switch's default baud rate is 9600.
- Verify that the switch is set up do PPP over the serial port. You can verify this by checking the settings in the Web interface (System→Settings).
- Verify that you clicked Reset from the Web interface when making changes to the screen (System→System Settings). This guarantees the serial port resets and initializes the modem. This is especially true with a modem connected to a switch that was restarted.
- Check the event log for failures.
- Make sure you have the correct dial-up networking settings. Refer to the section, [“Setting up a Dial-Up Networking connection.”](#)
- Make sure you have the remote modem set to auto answer and that it is in smart mode so that it can respond to the AT command set.
- Verify that the auto detection did not fail, and that the switch is in serial menu mode.

Cause:

You were dialed in and managing the switch remotely using PPP and you changed the baud rate and applied it, but now you cannot manage the switch.

Action:

To manage the switch, disconnect the dial-up connection and attempt to reestablish it. This gives the modem a chance to renegotiate the baud rate with the switch.

Cause:

You are set up to use PPP but want to use the serial port for the serial menu.

Action:

Choose the serial port mode Serial Menu. Press OK using the Web management interface (System→System Settings) and restart the switch. A serial cable must be installed in place of the modem in order to use the Serial Menu. Remember to power off the switch when plugging in and unplugging the serial port connection; otherwise, you might damage system components.

Cause:

You are set up to use the Serial Menu but want to use the port for PPP.

Action:

You can change the serial port settings (System→System Settings) or the Serial Menu itself. For these changes to take effect, restart the switch. For the best results, connect the modem while the switch is turned off.

Cause:

You are using a dial-up serial PPP connection and you encounter repeated CRC errors.

Action:

Make sure that the modem that is connected to the switch has hardware flow control enabled.

PPP option settings

These settings describe the switch's behavior when negotiating Serial PPP.

For IP:

- IP Address negotiation is enabled.
- The switch needs the peer's IP address to make a connection.
- The peer should not suggest an IP address for the switch. The switch uses its management IP address.
- The switch rejects VJ compression.
- The switch rejects VJ connection ID compression.

For LCP:

- The switch does not initiate a connection.
- The switch accepts magic number negotiation.
- The switch rejects address control field compression.
- The switch rejects protocol field compression.
- The switch does not allow asynchronous character map to be negotiated.
- The switch accepts Maximum Receive Unit (MRU) requests.

For authentication:

- The switch does not authenticate itself to a peer with PAP upon request.
- The switch requires that peers perform PAP authentication using the administrator's login and password.
- The switch does not authenticate itself to a peer with CHAP upon request.
- The switch does not require that the peer authenticate itself with CHAP.

Appendix C

System Messages

System forwarding (syslog) enables you to forward information from your switch's system log to different host machines using the system logging daemon (syslogd).

This appendix provides a listing of possible syslog messages that the Contivity VPN Switch might write to a remote system. Each message is followed by a description and the recommended corrective action, if any.

Certificate messages

Error removing CA certificate file: xxx

Description: The switch can be manufactured with a trusted CA certificate for use by SSL. The temporary manufacturing file containing the certificate is removed the first time you boot the switch. This error message indicates that the switch is unable to remove the temporary certificate file. The error might be caused by a general problem with the local file system.

Action: Manually delete all files in the /system/cert/ca directory.

Installed new CA certificate from file: xxx

Description: The switch can be manufactured with trusted CA certificates for use by SSL. This informational message indicates a trusted SSL CA certificate was installed when the switch was manufactured.

Action: No action required.

tCert: Shutdown complete

Description: This informational message indicates that the task responsible for certificate maintenance has shut down. This is usually part of the normal system shutdown.

Action: No action required.

tCert: task creation failed

Description: The task responsible for X.509 certificate maintenance on the switch failed to start properly. This most likely indicates severe resource exhaustion on the switch.

Action: Reboot the switch. If the reboot does not fix the problem, contact Nortel Networks Technical Support.

tCert: X.509 certificates disabled in flash memory

Description: This is an informational message that indicates the use of X.509 certificates by the switch has been totally disabled.

Action: No action required.

Warning: System CA certificates may have been tampered with, please reinstall!

Description: The switch performs a periodic integrity check of the SSL-related X.509 certificates that are stored on the switch's local file system. This message signals a failure during the integrity check. This indicates that one or more of the SSL-related certificates might have been tampered with, or that a certificate has been corrupted.

Action:

- 1 Delete, then reinstall any SSL-related certificates. It is not necessary to delete and reinstall the tunnel-related certificates since they are stored in the LDAP database and not on the local file system.

- 2 Manually verify the tunnel-related certificate fingerprints. You should perform this procedure any time you suspect tampering.

ISAKMP messages

ISAKMP [13] No proposal chosen in message from xxx (a.b.c.d)

In many cases, a Session:IPsec message precedes the ISAKMP message. If the Session:IPsec message indicates an error, then the Session message describes the cause and required action. If there is no Session:IPsec error message, refer to the following list of causes and solutions for explanations.

Description: The encryption types proposed by branch office xxx do not match the encryption types configured locally.

Action: Check the encryption types on both sides to make sure they match. If necessary, reconfigure the encryption on one system.

Description: The requested authentication method (for example, RSA Digital Signature) is not enabled.

Action: Enable all required authentication types. Make sure the unneeded types are disabled.

Description: One side of the connection is configured to support dynamic routing while the other side is configured for static routing, where branch office is xxx.

Action: Configure both sides to use the same routing type.

Description: Both sides are configured to support static routing. However, the local and remote network definitions of the two sides do not match, where branch office is xxx.

Action: Configure both sides to have matching local and remote network definitions.

Description: The Perfect Forward Secrecy (PFS) setting of the two sides do not match. Branch office xxx does not have PFS enabled, while PFS is required by the local settings.

Action: Make sure the PFS settings on both sides match. Either enable PFS on the remote side, or disable PFS locally.

ISAKMP [13] Error notification (No proposal chosen) received from xxx (a.b.c.d)

Description: The proposal made by the local switch has been rejected by an Extranet Access Client. This usually indicates that the client is using an international version (56-bit) while the switch has stronger encryption enabled.

Action: The encryption methods used by the client and the switch must match. Either provide the user with an Extranet Access Client version that supports the stronger encryption method used by the switch, or enable 56-bit encryption on the switch.

Description: The proposal made by the local switch has been rejected by a remote branch office switch, or by an IPsec implementation from another vendor.

Action: Check with the administrator of the remote system to determine the cause of the problem. If the remote system is another switch, the cause is noted in that system's log.

ISAKMP [13] Authentication failure in message from xxx (a.b.c.d)

In many cases, a Session:IPsec message precedes the ISAKMP message. If the Session:IPsec message indicates an error, the Session message describes the cause and required action. If there is no Session:IPsec error message, refer to the following list of causes and solutions for explanations.

Description: No encryption types are enabled for the account in question.

Action: Enable the desired encryption types.

Description: The requested authentication method (for example, RSA Digital Signature) is not enabled.

Action: Enable all required authentication types. Make sure the unneeded types are disabled.

ISAKMP [13] Error notification (Authentication failure) received from xxx (a.b.c.d)

Description: An Extranet Access Client attempted to connect, but the user supplied the wrong password.

Action: Make sure that the user and the switch have the same password.

Description: A remote branch office switch rejected your switch's attempt to authenticate.

Action: Contact the administrator of the remote system. If the remote system is another Contivity VPN Switch, the cause is noted in that system's log.

No response from client - logging out

Description: Your switch has lost network connectivity with the remote side.

Action: Verify the network connectivity between your switch and the remote side.

Description: A remote branch office using pre-shared key authentication is using a key that is different from what is configured on the local switch. Because the two sides are using a different encryption key, your switch cannot decrypt the encrypted messages from the other side, and therefore drops the messages.

Action: Make sure that both systems are using the same pre-shared key.

ISAKMP [13] xxx (a.b.c.d) has exceeded idle timeout - logging out

Description: The remote system has been idle (meaning that no traffic has been sent) for the amount of time configured in the Idle Timeout parameter (Profiles→Groups→Connectivity).

Action: If the Idle Timeout value is too low, increase it. To disable idle timeouts entirely, set the Idle Timeout value to 00:00:00.

ISAKMP [13] Invalid ID information in message from xxx (a.b.c.d)

Description: One side of the connection is configured to support dynamic routing while the other side is configured for static routing. Branch office is xxx.

Action: Configure both sides to use the same routing type.

Description: Both sides are configured to support static routing, however the local and remote network definitions of the two sides do not match. Branch office is xxx.

Action: Configure both sides to have matching local and remote network definitions.

ISAKMP [13] Error notification (Invalid ID information) received from xxx (a.b.c.d)

Description: One side of the connection is configured to support dynamic routing while the other side is configured for static routing. Branch office is xxx.

Action: Configure both sides to use the same routing type.

Description: Both sides are configured to support static routing. However, the local and remote network definitions of the two sides do not match. Branch office is xxx.

Action: Configure both sides to have matching local and remote network definitions.

Branch office messages

Couldn't install route for *remxxx@xxx*

Description: The switch was unable to install the route for the remote network (indicated by *remxxx@xxx*). This might result when the route collides with an existing static route.

Action: Remove the existing static route or change the route for the remote network to be a subset or superset of the static route.

SSL messages

Checking chain: invalid parent cert, xxx

Description: The given certificate in the chain is not valid. This might indicate that the certificate installed at the external LDAP server has expired or is invalid in some other way.

Action: Verify that the certificate is valid or use a certificate that you know is valid.

Checking chain: invalid child cert, xxx

Description: The given certificate in the chain is not valid. This might indicate that the certificate installed at the external LDAP server has expired or is invalid in some other way.

Action: Verify that the certificate is valid or use a certificate that you know is valid.

Child cert [xxx] not valid signature by [xxx] - xxx

Description: The given certificate in the chain is not properly signed. This error can indicate that the certificate was incorrectly installed at the external LDAP server.

Action: Reinstall the certificate at the external LDAP server.

Invalid root cert, xxx

Description: One of the root certificates passed to the switch during SSL negotiations was invalid.

Action: Configure the remote side to pass a valid chain of certificates to the switch.

No matching trusted CA certs

Description: None of the certificates in the chain are trusted CA certificates. This message can result if the CA certificate has not been installed or has not been marked as trusted on the switch.

Action: Make sure the CA certificate has been installed and that the certificate is marked as trusted on your switch.

Database messages

Configuration file: xxx does not exist

Description: The slapd.cnf file does not exist on the disk, therefore the internal LDAP server could not start. This error can occur if the switch disk has been modified.

Action: Reinstall the switch software.

Failed to start

Description: The internal LDAP server did not start. This can be caused by a missing configuration file.

Action: Reinstall the switch software.

Index file for attribute xxx from file xxx could not be created

Description: The given attribute index file for the internal LDAP server could not be created. This might indicate that the switch disk is full or that the database index files are corrupt.

Action: Restore the switch software from an FTP backup or reimport the database from the LDIF file.

LDIF file: xxx could not back up

Description: The internal LDAP server database could not be backed up to the specified LDIF file. This can result if the name of the LDIF file is not in 8.3 format.

Action: Make sure the backup file has an 8.3 file name.

LDIF file: could not restore xxx

Description: The internal LDAP server database could not be restored from the specified LDIF file. This might indicate that the LDIF file does not exist.

Action: Choose an LDIF file that currently resides on the switch disk.

Security messages

Account: xxx[xxx] uid xxx not found in account

Description: A UID of the remote entity was not found in the account used to initiate a branch office connection (the UID entry in the message is a UID for PPTP or L2TP, and a remote gateway address for IPsec). This error can result if the credentials given by the remote side of the branch office connection do not match the local configuration.

Action: Make sure the Remote Identity information of the IPsec Authentication Certificates section (Profiles→Branch Office→Edit Connection) is configured properly.

AuthServer: ldap inconsistent; no server type in entry xxx

Description: An LDAP entry for an authentication server does not contain a server type. This can indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

CaAuthServer: failed remove - xxx

Description: An LDAP entry for a CA authentication server was not fully created and then could not be removed. This can result if the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

CaAuthServerCollection: authenticate xxx cert [xxx] invalid signature by [xxx] - xxx

Description: The certificate passed in with the authentication request does not have a valid signature, based on the CA certificate configured on the switch. This can indicate either an incorrect certificate at the remote side (either a client or branch office), or an incorrect CA certificate was installed on the switch.

Action: Make sure that both sides have the correct certificates installed.

CaAuthServerCollection: authenticate xxx[xxx]:xxx bad certificate - xxx

Description: The certificate passed in with the authentication request is not a valid X.509 certificate. This error can result if the certificate configured either at the client or the other side of the Branch Office is incorrect.

Action: Install the correct certificates.

Security: store new system IP address xxx failed - xxx

Description: The system IP address could not be stored in the switch configuration LDAP entry. Possible cause: the LDAP server is not accessible.

Action: Start the LDAP server or change the external LDAP server configuration to make it accessible.

Security: store new system name xxx failed - xxx

Description: The system name could not be stored in the switch configuration LDAP entry. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Security: store new system subnet mask xxx failed - xxx

Description: The system subnet mask could not be stored in the switch configuration LDAP entry. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Entry is referenced [xxx] - xxx

Description: The LDAP entry is being referenced by another LDAP entry (for example, a filter set being referenced by a User Group or Branch Office Connection).

Action: Remove all references to the LDAP entry in question, then delete the entry.

Error copying entry [xxx] to [xxx] - xxx

Description: An error occurred while copying an LDAP entry.

Action: Delete the new copy that caused the error and retry the rename operation.

Error copying subentries of [xxx] to [xxx] - xxx

Description: An error occurred while copying a set of LDAP entries. This can be caused by an unreachable LDAP server.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Error copying tree [xxx] to [xxx] - xxx

Description: An error occurred while copying a tree of LDAP entries. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Error deleting entry [xxx] - xxx

Description: An error occurred while deleting an LDAP entry. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Error deleting tree [xxx] - xxx

Description: An error occurred while deleting a tree of LDAP entries. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

LocalAuthServer: failed remove - xxx

Description: An LDAP entry for an LDAP authentication server was not fully created and then could not be removed. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

SchemaCIs: Database schema not available

Description: The external LDAP server does not support a schema entry so it is not possible to update its schema over the network. This error can occur if the external LDAP server is not a Netscape server and does not support the cn=schema entry.

Action: Update the external LDAP server schema manually, then reconnect to it.

xxx xxx being referenced by xxx

Description: The LDAP entry is being referenced by another LDAP entry (for example, a filter set being referenced by a User Group or Branch Office Connection).

Action: Remove all references to the LDAP entry in question, then delete the entry.

Session: xxx uid invalid - authentication failed

Description: The given IPsec hashed UID was not found in the LDAP database. This can be caused if the UID typed in at the client was invalid or the account no longer exists.

Action: Make sure the correct UID was typed at the client and make sure the account is valid.

Session: xxx[xxx] invalid uid - authentication failed

Description: The given group UID was not found in the LDAP database, or the UID was found under a group account and this was not a group login. This error might result if the UID was mistyped at the client or the account no longer exists.

Action: Make sure the correct UID was typed at the client and make sure the account is valid.

Session: xxx[xxx] session rejected - system is initializing

Description: The switch had to reject the incoming request since it is still initializing.

Action: Wait a short time to make sure that the switch has initialized, then try again.

Session: xxx[xxx] session rejected - system is shutting down

Description: The switch rejected the incoming request since it is shutting down.

Action: Wait for the switch to restart, then try again.

Session: xxx[xxx]:xxx xxx auth method not allowed

Description: The authentication method of the incoming request is not allowed in the group that the session is bound to. The session is bound to a group by one of the following:

- The group that the user's account is in (in LDAP)
- The RADIUS default group
- The RADIUS class attribute
- The CA authentication server's default group

Action: Enable the authentication method for the bound group.

Session: xxx[xxx]:xxx - authentication failed using all authservers

Description: The incoming request could not be authenticated by any configured authentication servers (LDAP, RADIUS, or CA).

Action: Provide the correct credentials. (For example, create a new user account.)

Session: xxx[xxx]:xxx AddLink failed [xxx] current links xxx

Description: The multilink session could not be created. This can be caused by any of the following:

- New logins are disabled.
- The max sessions on the switch has been reached.
- There is not enough heap on the switch.
- The call admission priority slot is full.
- The call admission priority slot is outside of access hours.
- The max links configured for the group has been reached.

Action: Verify the correct settings for each of the possible causes.

Session: xxx[xxx]:xxx IP address assignment failed

Description: An address could not be assigned to the session. This can result if the static address for the session is in use or if the address pool is exhausted.

Action: Expand the number of addresses in the pool, or change the static address on the account.

Session: xxx[xxx]:xxx L2TP host [xxx] account misconfigured

Description: The L2TP Access Concentrator on the Branch Office Connection does not exist or does not have a LAC or switch UID.

Action: Recreate the L2TP Access Concentrator entry and make sure this entry is linked to the Branch Office Connection.

Session: xxx[xxx]:xxx account has max links (xxx)

Description: The maximum number of multilink sessions has been reached.

Action: Increase the maximum number of allowed PPP links on the Profiles→Groups→Edit→Connectivity screen.

Session: xxx[xxx]:xxx account has max sessions (xxx)

Description: The maximum number of sessions for the given account has been reached.

Action: Increase the number of logins on the Profiles→Groups→Edit→Connectivity screen.

Session: xxx[xxx]:xxx account is disabled

Description: The account is not currently enabled. This error can occur if the Branch Office Connection request is a different tunnel type than the local switch.

Action: Make sure that both sides are configured to support the same tunnel type.

Session: xxx[xxx]:xxx account not allowed now

Description: The session request is outside the permitted hours of access.

Action: Change the Access Hours setting assigned to the group on the Profiles→Groups→Edit→Connectivity screen.

Session: xxx[xxx]:xxx authentication failed using xxx

Description: The credentials for the session could not be validated by any of the authentication servers.

Action:

- 1 Make sure you are using the correct credentials.
- 2 Expand the capability of the RADIUS authentication server to handle the authentication method.
- 3 Add a new account with the given credentials.

Session: xxx[xxx]:xxx client assigned address [xxx] already in use

Description: The address given by the tunnel client is currently in use. This might indicate that the address is either being used in a static or dynamic route, or that the address is assigned to an active tunnel.

Action: Configure the client to use a different address.

Session: xxx[xxx]:xxx connect Qos level xxx full

Description: The switch does not have any more slots for the session's call admission priority. This can indicate that the configured Call Admission Priority for the group that the request is assigned to is too low.

Action: Increase the Call Admission Priority on the Profiles→Groups→Edit→Connectivity screen.

Session: xxx[xxx]:xxx invalid password - master admin authentication failed

Description: The primary administrator password was invalid. This can result from using the wrong password or from making a mistake while typing the password.

Action: Make sure you are using the correct password, and make sure you typed it correctly.

Session: xxx[xxx]:xxx login rejected - new logins disabled

Description: New logins are currently disabled. This can result if the switch was shut down with one of the following settings enabled on the Admin→Shutdown screen:

- The **Disable new logins** checkbox is selected, or
- The **Disable logins after restart** checkbox is selected.

Action: Enable new logins by deselecting the disable login settings on the Admin→Shutdown screen and then restart the switch.

Session: xxx[xxx]:xxx no memory free: xxx threshold: xxx

Description: There is not enough heap memory available to establish the session. This can result if the switch has consumed a large amount of memory while processing management requests.

Action: Increase the amount of physical memory on the switch, or wait until the management requests are complete.

Session: xxx[xxx]:xxx only one session/static address allowed

Description: An address can be used by only one session. This error occurs if the switch receives a second login to an account that has a static address configured.

Action: Change the account to use dynamic addresses from either a static address pool or DHCP.

Session: xxx[xxx]:xxx pool address [xxx] already in use

Description: The returned static pool address is currently in use. This error can occur if another tunnel is using this address via a static address configuration or another address pool. The error also occurs if a static host route using this address has been added.

Action: No action is necessary. The switch attempts to allocate a different address.

Session: xxx[xxx]:xxx session directed to use server xxx

Description: This is an informational message indicating that load balancing is enabled and the session is being redirected to another switch. This occurs when the switch is either more heavily CPU-loaded or session-loaded than the other switch.

Action: No action is necessary.

Session: xxx[xxx]:xxx static address [xxx] already in use

Description: The static address assigned to the account is in use by another tunnel or via a static host route.

Action: Change the static address.

Session: xxx[xxx]:xxx system has max sessions (xxx)

Description: The switch has reached its maximum number of sessions. This occurs when the switch reaches the maximum number of tunnels that can be configured.

Action: Use load balancing with another switch (if you are using IPsec clients), or upgrade the switch to the next higher model.

RADIUS accounting messages

RADIUS: Cannot send accounting request to <server-name>, possibly due to DNS translation failure

Description: This message indicates a connection failure. While sending a request, an error occurred due to a socket creation problem. This usually indicates a DNS resolution problem.

Action: Verify the following:

- DNS host name is correct.
- DNS server is configured properly.
- DNS server is available.

RADIUS: no reply from server <server-name>(<port number>)

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

RADIUS: <server-name> server timed out

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

RADIUS: network socket failure with <server-name>, recvfrom error: <error>

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

RADIUS: <server-name> server failed

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Indicated packet length too large

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

RADIUS: failure sending <user-name> accounting record to <server-name>

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Non-matching ID in server response

Description: This message indicates that an invalid response was received. The Transaction ID in the response packet is not the expected value.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Unsupported response type (<number>) received from server

Description: This message indicates that an invalid response was received. The response packet type is not one of the expected types: Access-Accept, Access-Reject, or Access-Challenge.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Received bad attribute type from server

Description: This message indicates that an invalid response was received. The RADIUS Attribute value is incorrect.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Response OK

Description: This message indicates that a valid response was received.

Action: No action necessary.

RADIUS: <user-name> accounting record sent to <server-name> OK

Description: This message indicates that a valid response was received.

Action: No action necessary.

RADIUS authentication messages

RADIUS: Cannot send request to <server-name>, possibly due to DNS translation failure

Description: This message indicates a connection failure. While sending a request, an error occurred due to a socket creation problem. This usually indicates a DNS resolution problem.

Action: Verify the following:

- DNS host name is correct.
- DNS server is configured properly.
- DNS server is available.

Login failure due to: Server network connection failure

Description: This message is received by the Extranet Access Client, and indicates a connection failure. While sending a request, an error occurred due to a socket creation problem. This usually indicates a DNS resolution problem.

Action: Verify the following:

- DNS host name is correct.
- DNS server is configured properly.
- DNS server is available.

RADIUS: no reply from RADIUS server <server-name>(<port number>)

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

**RADIUS: <server-name> server timed out authenticating
<user-name>**

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

**RADIUS: network socket failure with <server-name>, recvfrom
error: <error>**

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

**RADIUS: <server-name> server error while authenticating
<user-name>**

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Indicated packet length too large

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

RADIUS: <server-name> sent invalid response packet for <user-name>

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Non-matching id in server response

Description: This message indicates that an invalid response was received. The Transaction ID in the response packet is not the expected value.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Unsupported response type (<number>) received from server

Description: This message indicates that an invalid response was received. The response packet type is not one of the expected types: Access-Accept, Access-Reject, or Access-Challenge.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Received bad attribute type from server

Description: This message indicates that an invalid response was received. The RADIUS Attribute value is incorrect.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Invalid reply digest from server, possible shared secret mismatch

Description: This message indicates that an invalid response was received. The computed authenticator does not match the value in the packet.

Action: Verify that the shared secrets match.

RADIUS: <server-name> sent packet with invalid response authenticator for <user-name>

Description: This message indicates that an invalid response was received. The computed authenticator does not match the value in the packet.

Action: Verify that the shared secrets match.

RADIUS server returned access challenge

Description: This message indicates that a valid access-challenge response was received.

Action: No action required.

RADIUS: <server-name> sent challenge for <user-name>

A valid access-challenge response was received.

Action: No action required.

RADIUS access challenge received

Description: This message is received by the Extranet Access Client. A valid access-challenge response was received.

Action: No action required.

RADIUS server rejected access

Description: This message indicates that a valid access-reject response was received.

Action: No action required.

**RADIUS: <user-name> access DENIED by server
<server-name>**

Description: This message indicates that a valid access-reject response was received.

Action: No action required.

Response OK

Description: This message indicates that a valid access-accept response was received.

Action: No action required.

RADIUS: <user-name> access OK by server <server-name>

Description: This message indicates that a valid access-accept response was received.

Action: No action required.

Routing messages

Unable to create xxx for OSPF

Description: The switch could not create the necessary components to initialize OSPF. This could happen if the switch runs out of free memory.

Action: Disable and enable OSPF globally in Routing→OSPF screen. If this does not work, disable OSPF, boot the switch and enable OSPF in Routing→OSPF screen.

OSPF Disabled

Description: The administrator disabled OSPF from the Routing→OSPF screen.

Action: No action required.

Closing OSPF-RTM connection

Description: OSPF closed the RTM connection, which occurs if the administrator disables OSPF from Routing→OSPF screen.

Action: No action required.

Ospf_Global.State changed from ENABLED to DISABLED by user 'admin' @ x.x.x.x

Description: The administrator disabled OSPF from the Routing→OSPF screen.

Action: No action required.

Opened OSPF-RTM connection

Description: The administrator enabled OSPF from the Routing→OSPF screen and successfully registered with RTM.

Action: No action required.

OSPF Enabled

Description: The administrator enabled OSPF from the Routing→OSPF screen.

Action: No action required.

Ospf_Global.State changed from DISABLED to Enabled by user 'admin' @ x.x.x.x

Description: The administrator disabled OSPF from the Routing→OSPF screen.

Action: No action required.

Can not accept x.x.x.x as router id

Description: OSPF can not accept the given router ID in the Routing→OSPF screen.

Action: You must change router ID in the Routing→OSPF screen. Invalid router IDs are 127.0.0.1 and 0.0.0.0.

LoadOspfAreas Failed

Description: OSPF failed to load all areas of information from the config file. This could happen if the config file is damaged.

Action: Delete all OSPF areas, re-create them from the Routing→OSPF screen, and reboot the switch.

LoadOspfIntf Failed

Description: OSPF failed to load information for all interfaces from the config file. This could happen if the config file is damaged.

Action: Delete all OSPF interfaces, re-create them from the Routing→Interface screen, and reboot the switch.

VR xxx: Starting xxx as Master for xxx

Description: Logged when VRRP is starting as a master for an address. The parameters are:

- The VRID of this VR
- The reason for starting, either because it was enabled or the interface went up
- The IP address

Action: No action required.

VR xxx: Starting xxx as Backup for xxx

Description: Logged when starting as a backup for an address. The parameters are:

- The VRID of this VR
- The reason for starting, either because it was enabled or the interface went up
- The IP address

Action: No action required.

VR xxx: Starting xxx as master delayed Backup for xxx

Description: Logged when master delay mode is in effect. The parameters are:

- The VRID of this VR
- The reason for starting, either because it was enabled or the interface came up
- The IP address

Action: No action required.

VR xxx: Shutting down xxx on xxx

Description: Logged when VRRP is stopping. The parameters are:

- The VRID of this VR
- The reason for stopping, either because it was disabled or the circuit went down
- The IP address

Action: No action required.

Unable to get configuration for VR xxx

Description: This is an error event. It is logged when VRRP is enabled but the common configuration parameters are missing. These are the items set under the Routing→VRRP screen. The parameter is the IP address that is missing information.

Action: No action required.

RIP xxx: RIP Enabled

Description: Logged when RIP is globally enabled.

Action: No action required.

RIP xxx: RIP Disabled

Description: Logged when RIP is globally disabled.

Action: No action required.

RIP xxx: Can't alloc main node

Description: Logged when there is not enough memory to allocate RIP parameters.

Action: No action required.

RIP xxx: Circuit xxx created

Description: Logged when the RIP circuit was created. The parameter stands for circuit ID.

Action: No action required.

RIP xxx: Circuit xxx deleted

Description: Logged when the RIP circuit was deleted. The parameter stands for circuit ID.

Action: No action required.

RIP xxx: Unable to register with UDP

Description: Logged when you cannot register with UDP protocol.

Action: No action required.

RIP xxx: setsockopt RIP socket xxx SO_RCVBUF xxx failed

Description: Logged when RIP receive buffers are not large enough. This can happen when a large numbers of RIP neighbors send their RIP updates simultaneously. The first parameter is the socket number and the second parameter is the maximum receive buffer size.

Action: No action required.

RIP xxx: bind RIP socket xxx failed

Description: Logged when RIP failed to bind the socket.

Action: No action required.

RIP xxx: Unable to spawn Dispatcher task xxx for RIP

Description: Logged when RIP failed to spawn the main task responsible for receiving RIP packets. The parameter stands for the name of the task.

Action: No action required.

RIP xxx: Unable to spawn timer task xxx for RIP

Description: Logged when RIP failed to spawn the timer task. The parameter stands for the name of the task.

Action: No action required.

RIP xxx: cid xxx mismatched auth password from xxx

Description: Logged when RIP authentication failed while receiving RIP packets. The first parameter is the circuit ID on which it was receiving RIP packets and the second parameter is the IP address from which it received RIP packets.

Action: No action required.

Hardware messages

The switch software provides informational messages when cards are removed and replaced. When you exchange two cards with each other, the switch considers this two simultaneous replacements.

Interface [nnn] not present, deleting from config

Description: This indicates that the configuration file contains an interface [nnn] entry, but there is no card in the slot. The interface [nnn] entry is deleted from the configuration.

Action: No action required.

Interface [nnn] replaced, resetting config

Description: This indicates the card type specified in the configuration file does not match the card type currently in the slot. The configuration information is reset to defaults then initialized with the current hardware.

Action: No action required.

Interface [nnn] replaced, deleting from config

Description: This indicates the card type specified in the configuration file does not match the card currently in the slot. The interface is deleted from the configuration. This applies when the replaced card has more ports than the current card.

Action: No action required.

HWAccel [nnn] not present, deleting from config

Description: This indicates the configuration file contains a HWAccel [nnn] entry, but there is no hardware accelerator in the slot. The HWAccel [nnn] entry is deleted from the configuration.

Action: No action required.

Glossary

acknowledgement (ACK)

A type of message sent to indicate that a block of data arrived at its destination without error.

address masks

IP addresses used to represent a series or range of IP addresses.

authentication

A security procedure where a user verifies his identity before accessing networks protected by a firewall.

bandwidth

The difference between the highest and lowest frequencies of a transmission channel; amount of data that can be sent through a given communications circuit.

certification authority (CA)

An authority that issues digital certificates and manages the life cycle of certificates.

Challenge Handshake Authentication Protocol (CHAP)

A peer entity authentication method for PPP, using a randomly-generated challenge and requiring a matching response that depends on a cryptographic hash of the challenge and a secret key.

client

A system or process that requests a service of another system or process.

default route

A route that is used when the switch receives traffic for which no matching route is in the routing table.

Diffie-Helman

A key agreement algorithm that does key establishment, not encryption. However, the key it produces may be used for encryption, for further key management, or any other cryptography.

digital certificate

A certificate document in the form of a digital data object to which is appended a computed digital signature value that depends on the data object.

distinguished name (DN)

An identifier that uniquely represents an object in the X.500 Directory Information Tree. An X.509 public-key certificate or CRL contains a DN that identifies its issuer, and attribute certificate identifies its subject.

Domain Name System (DNS)

A general purpose distributed, replicated, data query service used to look up host IP addresses based on host names. DNS applications can perform name-to-address and address-to-name translations.

dynamic routes

Routes that are learned via the switch's RIP support, and are used for branch office connections and the private interface.

encryption

The manipulation of a packet's data to prevent any but the intended recipient from reading the data.

encryption certificate

A public-key certificate that contains a public key that is intended to be used for encrypting data, rather than for verifying digital signatures or performing other cryptographic functions.

Federal Information Processing Standards (FIPS)

Technical guidelines for U.S. Government procurements of information processing system equipment and services.

File Transfer Protocol (FTP)

A TCP-based application layer Internet Standard protocol that transfers files to and from a remote host.

firewall

A collection of hardware and software components that controls communication between two networks, such as a private network and the Internet. All information passed between the two networks must pass through the firewall. The firewall allows only authorized traffic to pass between the networks.

gateway

A communications device or program that passes data between networks having similar functions but dissimilar implementations.

integrated firewall

The switch supports two integrated firewalls: the Contivity Firewall and the FireWall-1 Inspection Module from Check Point Technologies, LTD. Also referred to as an embedded firewall.

interface

The connection between a router and one of its attached networks. An interface to a network has a single IP address and mask associated with it.

Internet

The single, interconnected, worldwide system of commercial, government, educational, and other computer networks that share a set of protocols.

Internet Protocol (IP)

The transport layer protocol used by the Internet Protocol family for transporting information among computers.

Internet Security Association and Key Management Protocol (ISAKMP)

Defines how encryption keys for sessions are initiated and updated.

intranet

A computer network, especially one based on Internet technology, that an organization uses for its own internal, and usually private, purposes and that is closed to outsiders.

IP address

The identifiers used by the protocols that govern Internet information exchange. The Internet Network Information Center assigns these numbers to uniquely identify different machines on the Internet.

IPsec

A tunneling protocol that offers a strong level of encryption, integrity protection. It uses digital certificates, password-based keys, and tokens for authentication.

IPsec Key Exchange (IKE)

An Internet IPsec key-establishment protocol that puts in place authenticated keying information for use with ISAKMP and for other security associations.

key agreement

A method for negotiating a key value without transferring the key, even in an encrypted form, such as Diffie-Helman.

Layer2 Tunneling Protocol (L2TP)

Tunneling protocol that enables secure remote access to enterprise networks across the Internet.

Lightweight Directory Access Protocol (LDAP)

Protocol based on directory entries that provide access for management and browser applications that provide read/write interactive access to the X.500 directory.

local area network (LAN)

A data network intended to serve an area of only a small area to optimize data transfer rates.

management information base (MIB)

The set of parameters an SNMP management station can query or set in the SNMP agent of a network device, such as a router.

management IP address

The IP address that is used to manage all system services from a Web browser, such as HTTP, FTP, and SNMP. This address must be accessible from one of the switch's private physical interfaces. To be accessible, the Management IP Address must map to the same network as one of the private interfaces.

management station

The remote workstation that is used to configure and manage the Contivity VPN Switch's integrated Check Point firewall. The FireWall-1 Management Console software runs on Management Station. The GUI portion of the Management software can be installed on the Management Station or on a different system.

medial access control (MAC) address

The hardware address of a device connected to a shared media.

Network Address Translation (NAT)

A mechanism that converts an internal network's private addressing scheme to an acceptable Internet address, thereby enabling the internal systems to communicate on the Internet.

Network Time Protocol (NTP)

Synchronizes the clocks of various devices across networks.

Open Shortest Path First (OSPF)

OSPF is a link-state routing protocol that maintains a database from which a routing table is constructed from the shortest path, using a minimum of routing protocol traffic.

packet

The unit of data sent across a network. Typically, it refers to application data units.

PING

A program used to test reachability of destinations by sending an ICMP echo request and waiting for a reply.

Point-to-Point Protocol (PPP)

A protocol that provides a method for transmitting packets over serial point-to-point links.

Point-to-Point Tunneling Protocol (PPTP)

A tunneling protocol that is used as a security tool.

port

A transport layer demultiplexing value. Each application has a unique port number associated with it.

private default route

The default routes that are used for traffic that comes into the switch via a public interface, through a tunnel, or from the switch's public interface address.

protocol

A formal description of message formats and the rules two computers must follow to exchange those messages. A protocol can describe low-level details of machine-to-machine interfaces or high-level exchanges between application programs.

public default route

The default routes that are used for traffic that comes into the switch via a private interface or from the switch's private interface address.

Resource Reservation Protocol (RSVP)

A protocol used to signal QoS requests and confirmations.

route

The path that network traffic takes from its source to its destination.

router

A device that forwards traffic between networks. The forwarding decision is based on network layer information and routing tables.

routing

The process of selecting the correct interface and next hop for a packet being forwarded.

Routing Information Protocol (RIP)

A distance vector, as opposed to link state, routing protocol.

RSA digital signature

A public-key cryptographic system that may be used for encryption and authentication.

server

A provider of resources, such as file servers and name servers.

Simple Network Management Protocol (SNMP)

The Internet standard protocol developed to manage nodes on an IP network.

split horizon

A method that RIP uses to avoid routing problems caused by including routes in updates sent back to the gateway from which they were learned. The simple split horizon scheme omits routes learned from one neighbor in updates sent back to that neighbor. An extension to this method is called split horizon with poisoned reverse. It includes the learned routes, but assigns them a cost of infinity, which causes an update.

static routes

Routes that are manually configured in the switch's routing table.

stub network

A network that only carries packets to and from local hosts. Even if it has paths to more than one other network, it does not carry traffic for other networks.

subnet

A portion of a network, which may be a physically independent network segment, that shares a network address with other portions of the network and is distinguished by a subnet number.

Telnet

A command protocol used to establish login sessions on a remote host.

triggered update

A method used by RIP in which a new routing table is sent almost immediately after a routing change has been made. This is in contrast to the poison reverse method, in which routes are updated after a cost of infinity is reached, a process that can take much time.

User Datagram Protocol (UDP)

An Internet standard transport layer protocol. It is a connectionless protocol that adds a level of reliability to an multiplexing to IP.

Uniform Resource Locator (URL)

A standard for identifying objects on the Internet accessible through the Web.

Virtual Router Redundancy Protocol (VRRP)

A protocol that handles private interface failures. VRRP targets hosts that are configured with static next-hop routing addresses or default gateways. It provides a means of rerouting traffic in the event of a system/interface failure.

Wide Area Network (WAN)

A network, usually constructed with serial lines, that covers a large geographic area.

Index

A

- access control filtering 53
- access hours 35, 48, 120
- accessible networks 140, 147
- accessing the Contivity VPN Switch 68
- accounting
 - data 80
 - overview 59
 - records 79, 81
- accounting log 286
- ACE 199
- active sessions 318
- ActiveX Scripts 314
- administrator 92
 - settings 271
- authentication
 - Allow All 224
 - details 194
 - failed 295
 - group password 199
 - individual instances 224
 - overview 49
- authentication methods 35, 48, 120
- authentication servers
 - configuring 193
- authorization code 228
 - Entrust 235
- automatic backup 273
- AXENT tokens 204

B

- background images 317
- backup
 - configuration 277
 - host server(s) 272
- backup and restore 83
- bandwidth guarantees 56
- bandwidth management 55, 82
- base group 112
- blanket authentication 224
- branch office
 - access hours 35, 48, 120
 - accessible networks 140
 - authentication methods 35, 48, 120
 - configuring 119
 - encryption settings 122
 - indirectly connected switches 121
 - sample configurations 120
 - sample procedure 129
- browser error messages 316
- browsing delays 315
- bulkload 69

C

- CA certificate 226
 - lifetimes 221
 - obtaining from PKI 231
- cache 316
- call admission
 - guarantees 57
 - priority 57, 115

- cards
 - LAN 87
 - WAN 87
 - certificate
 - expiration 223
 - maintenance 222, 237
 - termination 223
 - certificate authority (CA) 220
 - firewalls 223
 - network controls 223
 - certificates
 - Microsoft Certificate Authority 237
 - cestraps.mib 327
 - CHAP 47
 - RADIUS 200
 - Check Point Firewall-1 39, 187
 - children of base group 112
 - class attributes RADIUS 202
 - client 197
 - client address redistribution 183
 - sample 184
 - summarization 184
 - client policy 155
 - setting up 72
 - color setting 318
 - Command Line Interface (CLI) 62, 69
 - compression 47
 - configuration
 - choices 99
 - IP address 88
 - log 283
 - preparing for 97
 - Web interface options 98
 - configuration file 76
 - local 61
 - storing 59
 - configuration log 287
 - configuration tasks
 - initial 69
 - ongoing 78
 - configuring
 - branch offices 119
 - serial interface 91
 - connectivity problems
 - overview 289
 - solving 292
 - control tunnels 69, 140
 - branch office 70, 142
 - creating 144
 - nailed-up 74, 144
 - restricted mode 73, 143
 - sample 143
 - types 70
 - user 71, 142
 - CRL 225
 - checking 232
 - servers 226
 - status 236
 - CSU/DSU 103
- ## D
- data collection records 80
 - data storage 79
 - default
 - gateway 90, 96
 - login 97
 - password 97
 - default route 176
 - branch office 121
 - Defender 199
 - DES 45
 - encryption 303
 - DHCP server 213, 304
 - dial-up
 - monitor 292
 - problems 293
 - Dial-Up Networking 311
 - Diffie-Hellman 197

- digital certificates
 - authenticating a client 236
 - components 219
 - maintaining 237
 - quick start 226
 - SSL 51, 215
- direct connections 33
- DNS 47, 95
 - server 312, 315
- docking station configurations 299
- domain controller 304

E

- encryption
 - overview 52
 - settings for branch office 122
- encryption methods 75
- enhanced routing 37, 166
- enterprise environment 66
- Entrust 52, 218
 - certificate store 222
 - PKI 220, 227
- Entrust Manager digital certificates 52, 218
- Entrust PKI 226
- Entrust PKI CA 223
- event log 283, 286
- External DHCP server 318
- extinction
 - interval 306
 - timeout 306
- Extranet Access 30
 - client monitor 290
 - connection problems 294

F

- factory default 275, 277
 - configuration 277
- filters through a firewall 35

- FIPS
 - overview 63
- firewall
 - branch office 120
 - imbedded 39, 186
 - integrated 39, 186
 - interaction with branch office 121
 - routing 168
 - statistics 286
- forwarding priority 56, 115
 - quality of service 56

G

- gateway IP address 92
- general problems
 - overview 290
 - solving 314
- generate certificate request 228
- group hierarchy
 - configuration 113
 - inheritance 113
- group password authentication 199
- guided configuration 53, 98
 - requirements 100

H

- hard drive, reformatting 278
- hardware
 - health check 285
- hardware encryption accelerator 107
- HDLC framing 301
- health check 79, 285
- health check display 318
- historical event logging 283
- HMAC 198
- HTTP 315
- HTTP management 92
- Hyperterminal 293

I

- import tunnel certificate 221
- indirectly connected switches 121
- inheritance, group 116
- inner IP address 213
- internal address pool 318
- Internet Explorer 317
- Internetwork Packet Exchange 165, 188
- IP accept policies 170
- IP address 90
 - assigning 41
 - configuration utility 88, 94
 - reconfiguring values 93
- IP address system 95
- IP announce policies 170
- ipconfig command 292
- IPSec 29, 44, 45, 197
 - password 295
 - tunnel filters 35
 - tunneling 44
 - username 295
- IPSecShm 299
- IPX 43, 188
- is 237
- ISAKMP 45, 52, 218
- ISAKMP/IPSec filters 35
- ISP 103
 - environment 66
 - tunnels 112

J

- Java 314, 317
- JavaScript 314
- jetpack.exe 307

K

- key pairs certificates 223

L

- L2F 29, 44, 47
- L2TP 29, 44
 - configuring over IPSec 155
- L2TP/IPSec client 160
- L2TP/IPSec gateway 161
- L2TP/L2F
 - filters 35
- last name search 119
- LCP options 302
- LDAP 75, 193, 204
 - attribute search 119
 - authentication 50
 - database authentication 198
 - database files 59
 - directory 205
 - full vendors 205
 - Netscape Directory Server 211
 - overview 50
- LDAP database
 - backup and restore 83
 - files 76
- license keys 88
- licenses 99
- Lightweight Directory Access Protocol 204
- logging 59, 62
 - displays 291
 - files 76
 - files, large 319
- login 97
 - ignored 317
 - not allowed 295
- logs
 - accounting 286
 - events 286
 - security 287
 - system 287
- loopback test 301

M

- management
 - Extranet Access Switch 100
 - IP address serial configuration 91
 - IP subnet mask serial configuration 91
 - overview 65
 - tasks 53
- manager WAN statistics display 301
- master browser 307
- maximum number of sessions 295
- MD5 45
- MIB support 323
- Microsoft
 - auto disconnect feature 297
 - client troubleshooting tools 292
 - Internet Explorer 314
 - Knowledge Base 313
 - networking tips 304
- modem hardware errors 303
- MS NetWare client 190
- MS-CHAP 47
- MS-DOS naming convention 319
- multiple Help windows 317

N

- NAT
 - dynamic 125
 - dynamic pooled 126
 - dynamic port 125
 - rules 129
 - sets 128
 - translation type 129
- NetBEUI 298, 305
- NetBIOS 298, 304, 305, 310
- NetMedic 292
- Netscape
 - Communicator 314
- Netscape Directory Server

- LDAP 206, 211
- netstats command 292
- network access server (NAS) 47
- Network Address Translation (NAT) 122
- network configurations 32
- Network Neighborhood 305
- network operations center (NOC) 65
- newoak.mib 328
- Normal mode 73
- Nortel Networks MIB 281
- Novell intraNetWare client 190

O

- Optivity VPN Manager 83
- OSPF overview 38, 179
- outer IP address 213

P

- PAP 47
 - RADIUS 200
- password 97
- performance problems
 - overview 290
 - solving 303
- ping command 294, 298
- PKCS #10 220
- PKI
 - adding remote clients 233
 - infrastructure 220
- poison reverse 179
- power failure 318
- PPP 63
 - layer 302
- PPTP 44, 46
 - control connection filters 35
 - tunnels filters 35
 - white papers 311

- Preside 84
 - applications 84
- primary administrator 92
- primary WINS server 299
- private key 223
- private LAN 40, 102
- public data network (PDN) 40
- public key 223
- public key infrastructure (PKI) 52, 218

Q

- quality of service
 - forwarding priority 56
 - overview 54
 - RSVP 58
- Quick Start 100
 - configuration 53

R

- RADIUS 75, 193, 200
 - accounting 80, 193
 - authentication 50, 198
 - class attributes 202
 - configuring server 202
 - external server 75
 - overview 51
- RC4-128 215
- RC4-40 215
- recovery diskette 275
- reference number 228
 - Entrust 220, 235
- register 99, 100
- relative distinguished name 115
- remote access 31
 - branch office 161
 - configuring 155
 - configuring the CES 155
 - configuring Windows 2000 160

- remote connectivity tunneling 44
- renewal interval 306
- reports 79, 285
- reservation control filters 35
- Reset button 279
- restart failure 320
- RIP 38, 177
 - RFCs 177
 - using 178, 179
- root CA 237
- route table 77
- routes
 - default 176
 - dynamic 37, 173
 - private 174
 - public 174
 - static 37, 173
- routing
 - enhanced 168
 - loops 178
 - overview 36, 165
 - route lookup 175
 - services 169
 - table 37, 173
 - traffic patterns 181
 - VPN 37, 166
- routing policy service 170
- routing table 37, 173
 - lookup 175
- RSA digital signature 226
 - authentication 232
 - certificates 52, 218, 222
 - enabling 232
- RSVP quality of service 58

S

- Safe mode 73, 93
- search for users 119
- Secure Socket Layer 215

- SecurID 197
 - security association 197
 - Security Dynamics 197
 - security log 283, 287
 - serial interface
 - configuration 91
 - prerequisites 91
 - serial number 277
 - serial number search 89
 - serial PPP
 - dial-up networking 342
 - establishing a connection 341
 - option settings 347
 - overview 341
 - setting up the switch 343
 - troubleshooting 345
 - server certificate 226, 227
 - importing 230
 - service level agreements 82
 - Service Pack 2 311
 - services 99
 - services routing 37, 166, 169
 - sessions 284
 - SHA 45
 - SHA-1 197
 - SMNP version 1
 - traps 281
 - SNMP
 - MIB II 54
 - overview 54
 - support 76
 - traps 54, 77, 291
 - software versions 275
 - split horizon 178
 - split tunneling
 - enable 149
 - third-party clients 151
 - split-horizon DNS 298
 - SSL digital certificates 51, 215
 - static address NAT 124
 - static routes 167
 - statistics 286
 - memory 79
 - statistics display 291
 - status 282
 - status reports 79
 - subject DN 136, 137, 156, 160, 224
 - subnet mask 90, 96
 - subnetworks 35, 48, 119
 - Switch concepts 29
 - synchronize RADIUS servers 205
 - system
 - log 283
 - shutdown 281
 - status 285
 - system identity 95
 - system log 287
 - system messages 349
 - branch office 354
 - certificate 349
 - database 356
 - hardware 380
 - ISAKMP 351
 - RADIUS accounting 367
 - RADIUS authentication 370
 - routing 374
 - security 357
 - SSL 355
- T**
- T-1 103
 - T1/V.35 interface 301
 - Telco expenses reduced 32
 - terminal emulator 91
 - third-party clients
 - authentication methods 151

- configuring 150
- split tunneling 151
- supported 150
- time zone 226, 233
- tokens
 - card 199
 - security 197
- tracert command 292
- traffic patterns 181
- translation types 129
- translation types, NAT 124
- traps
 - hardware 329
 - information for all 336, 337
 - intrusion-related 336
 - login-related 335
 - server-related 333
 - software-related 335
 - system-related 336
- triggered updates 179
- triple DES 45, 303
- troubleshooting
 - client address redistribution 320
 - client connection problems 294
 - Extranet Access Manager 314
 - Internet service provider login problems 313
 - modem and dial-up problems 293, 313
 - overview 289
 - PPTP connectivity 313
 - routing 320
 - toolbox 290
 - WAN link problems 300
- tunnels, configuring 111

U

- Uniform Resource Locator (URL) 97
- update frequency cpl 232
- upgrade fails 319
- upgrading software 315

- user
 - ID search 119
- user group
 - adding 74
- users
 - adding to a group 75
- Utunnel 183

V

- verify interval 306
- VeriSign OnSite configuration 221
- VPN connector
 - digital certificates 52, 218
 - Entrust 221
- VPN routing 37, 166
- VRRP overview 38, 180

W

- WAN interfaces
 - configuration 103
 - display 301
- Web browser 68
 - problems 314, 318
- Web Connector
 - digital certificates 52, 218
 - Entrust 221, 229, 231
- Web interface options 98
- Welcome display 98
- winipcfg command 292
- WINS 47
 - secondary servers 306
 - server 305, 310
 - settings 306
- Winsock DNS Update 312

X

- X.509 certificates 52, 218, 220, 222, 236